Predictive Power of Markets

- Prediction Accuracy, Incentive Schemes, and Traders' Biases -

Zur Erlangung des akademischen Grades eines
Doktors der Wirtschaftswissenschaften

(Dr. rer. pol.)

von der Fakultät für
Wirtschaftswissenschaften
der Universität Karlsruhe (TH)

genehmigte

DISSERTATION

von

Dipl.-Wirtsch.Inf. Stefan Luckner

Tag der mündlichen Prüfung: 25. Februar 2008

Referent: Prof. Dr. Christof Weinhardt

Korreferent: Prof. Dr. Rudi Studer

2008, Karlsruhe

Abstract

Accurate predictions are essential in many areas such as business and sports forecasting. Prediction markets are a promising approach for predicting uncertain future events and developments. To give a few examples, prediction markets have been employed successfully to aggregate information on the expected outcome of elections, sports events, and Oscar winners. This work studies the prediction accuracy of markets in the field of sports forecasting as well as the impact of traders' biases on their trading behavior. Traders indeed exhibit a substantial amount of biases in markets which were run for predicting the outcome of the FIFA World Cup 2006. Despite these biases, an empirical comparison of the markets and predictions derived from the FIFA world ranking, i.e. historic data, and betting odds shows that prediction markets are more accurate predictors than the FIFA world ranking and as accurate as betting odds from professional bookmakers. Betting odds, in turn, are known to predict extremely accurately. Traders' biases thus do not necessarily lead to poor predictions in case of prediction markets.

Another focus of this work is to study the impact of different monetary incentive schemes for play-money prediction markets on the accuracy of predictions. In order to do so, predictions from three groups of traders, corresponding to three treatments with different and widely-used incentive schemes, are compared with regard to their prediction accuracy in a field experiment. Subjects of the first group were paid a fixed amount, subjects of the second group were paid according to their ordinal rank within the group, and in case of the third group the subjects' payments depended linearly on their deposit value in the prediction market. The highest correlation between the relative frequency of outcome and trading prices is found in case of the second group, the rank-order tournament. Somewhat surprisingly, the rank-order tournament seems to beat the third incentive scheme where the traders' payments are based linearly on their return in the market.

Overall, this work demonstrates that markets are accurate predictors beyond the field of political stock markets. Moreover, the findings on traders' biases and incentive schemes are valuable for designing future prediction markets.

Table of Contents

List of Abbreviations

CA	Call auction
CDA	Continuous Double Auction
DPM	Dynamic pari-mutuel market
DV	Deposit value
ESA	Economic Science Association
EU	European Union
FIFA	Fédération Internationale de Football Association
FP	Fixed payment
GDP	Gross Domestic Product
HP	Hewlett-Packard
HSX	Hollywood Stock Exchange
IEM	Iowa Electronic Markets
MM	Market maker
MSR	Market scoring rule
NBA	National Basketball Association
NFL	National Football League
PSM	Political Stock Market
RO	Rank-order tournament
UBC	University of British Columbia
UK	United Kingdom

US	United States
USA	United States of America

List of Figures

List of Tables

1. Introduction

Uncertainty and doubt are seen to be major challenges for management in the 21st century (Nohria and Stewart, 2006). Considering the environment in which organizations are acting today, this is not surprising: Increasing speed of innovation and thus shorter product life cycles as well as the globalization of markets make our world increasingly complex and unpredictable. Hence, for organizations it is more important than ever to develop foresight capabilities to better foresee future developments, trends, potentials, challenges, and risks.

Predicting the future is an integral part of corporate decision making. Inaccurate or delayed predictions can result in substantial costs for a company. Improving foresight capabilities, on the other hand, helps to strengthen the position of a company in global competition. Most business challenges related to, for example, demand forecasting and new product development require information which is dispersed among many people. However, these people cannot be easily identified in most cases. But more and more companies recognize the potential of collective intelligence and try to leverage the *wisdom of crowds*[1] through technologies such as wikis, blogs, or reputational systems. All of these technologies help to aggregate information and gain a better understanding of the future by collecting knowledge of as many people as possible.

1.1. Motivation

Over the last couple of years, interest in prediction markets as a forecasting method has continuously increased in the scientific world and in industry. Markets provide incentives for information revelation and can be used as a mechanism for aggregating information. So far, prediction markets have done well in every known comparison with other forecasting methods (Hanson, 2006). Racetrack odds beat horse experts consistently (Figlewski, 1979), orange juice futures have proven more accurate than the National Weather Service of the US Department of Commerce (Roll, 1984), and stock prices determined the company responsible for the explosion of the Challenger spacecraft within 13 minutes – four months before a panel of experts published its

[1] Surowiecki (2004) created public interest in collective intelligence with his bestselling book "The Wisdom of Crowds".

official report (Maloney and Mulherin, 2003). Whereas information aggregation is only a byproduct of most traditional markets, prediction markets are set up with the explicit purpose of soliciting information. Engineering carefully, prediction markets can directly guide decision making.

The basic idea of prediction markets is to trade contracts whose payoff depends on the outcome of uncertain future events. Although the final payoffs of the contracts are unknown during the trading period, rational traders should sell contracts if they consider them to be overvalued and buy contracts if they consider them to be undervalued (Glosten and Milgrom, 1985). Until the outcome is finally known, the trading prices reflect the traders' aggregated beliefs about the likelihood of the future events. In informationally efficient markets, all the available information is reflected in the trading prices at any time (Fama, 1970a, Fama, 1991).

Examples of prediction markets that are open to the public include the Iowa Electronic Markets[2], the Political Stock Market PSM[3], TradeSports[4], NewsFutures[5], the Hollywood Stock Exchange[6], and STOCCER[7]. Several major companies such as Hewlett-Packard, Google, or Microsoft are also using internal prediction markets for company-specific predictions. The results of recent studies on these prediction markets are encouraging. One of the main reasons for their dissemination is that they have shown a high prediction accuracy compared to traditional forecasting methods such as polls, expert predictions, or surveys (Berg et al., 2001, Servan-Schreiber et al., 2004, Spann and Skiera, 2003). Good performance has also been demonstrated in corporate environments (Chen and Plott, 2002, Ortner, 2000, Plott, 2000). Beyond prediction accuracy, markets also provide considerable advantages in terms of continuous forecasting, participation, and cost efficiency compared to other widespread forecasting methods.

Continuous scanning of ongoing developments as an input to strategic planning may be difficult to implement with traditional forecasting methods such as brainstorming

[2] http://www.biz.uiowa.edu/iem
[3] http://psm.em.uni-karlsruhe.de
[4] http://www.tradesports.com
[5] http://us.newsfutures.com
[6] http://www.hsx.com
[7] http://www.stoccer.com

techniques, Delphi studies, and scenario workshops. The results of suchlike approaches usually have to be manually analyzed, evaluated, and summarized. All of this has to be performed at a certain point in time. In contrast, all the traders' information is aggregated by the price mechanism of a prediction market. This has two positive effects: First, the information aggregation by the price mechanism reduces the workload compared to traditional forecasting methods. Second, the price mechanism ensures that trading prices continuously reflect the totality of previously revealed knowledge and immediately respond to new information (Hanson, 1999). This means that information aggregated via prediction markets is available in the market and always up-to-date (Berg et al., 2003).

Concerning participation in foresight studies, it is a well-known problem that people generally refuse to participate or drop out early due to other commitments they consider more important (Cuhls, 2003). Therefore, it makes sense to provide incentives for participation. With proper incentive schemes traders do not necessarily state their individual preferences but their true beliefs (van Bruggen et al., 2006). Prediction markets allow for rather sophisticated incentive schemes as traders can be rewarded based on their performance, i.e. the quality of their contributions. This can happen in different ways. The market operator can for instance award prizes or money to the best traders or traders can be asked for investing some of their own money in a market. Yet, it is sometimes not even essential to provide monetary incentives or prizes to motivate participation. Prediction markets have also shown to perform well without providing any monetary incentives, e.g. by publicly announcing a ranking based on the traders' success in the market (Christiansen, 2007).

The implementation of a foresight activity is often restricted due to tight budget constraints and other resource limitations (Salo and Cuhls, 2003, Clar, 2003). As described above, the information aggregation process in prediction markets is carried out via the price mechanism and does not require any manual intervention. Prediction markets are highly scalable as the workload of the operators is almost independent from the number of traders and the time horizon (Chan et al., 2002). Furthermore, the hardware costs for running a market are negligible once the market platform has been designed and developed (Spann et al., 2007).

To sum up, evidence so far suggests that prediction markets are at least as accurate as traditional forecasting methods. Furthermore, they provide considerable advantages in terms of continuous forecasting, participation and information revelation as well as scalability and cost efficiency. This also explains why prediction markets currently receive a lot of attention in research.

This work studies the prediction accuracy of markets in the field of sports forecasting as well as the impact of traders' biases on their trading behavior. Data from predictions markets which were run for predicting the outcome of the FIFA World Cup 2006 is used to find out whether traders' biases lead to poor predictions. Furthermore, the markets' predictions are empirically compared to predictions derived from historic data and betting odds. Another focus of this work is to explore the impact of different monetary incentive schemes on the prediction accuracy of play-money markets. In order to do so, predictions from three groups of traders, corresponding to three treatments with different and widely-used incentive schemes, are compared with regard to their prediction accuracy in a field experiment.

1.2. Research Questions

The main objective of this work is to demonstrate the predictive power of markets in general and in the field of sports forecasting in particular. Moreover, the research on traders' biases and incentive schemes is valuable for designing future prediction markets. Within the scope of this work, the following research questions are addressed:

(I) How well do markets predict the future?

As was already mentioned prediction markets seem to outperform traditional forecasting methods in many cases. An evaluation of their prediction accuracy relative to traditional forecasting methods such as expert opinions or polls is required to answer the first research question. Earlier empirical research on prediction markets substantiates their predictive power in several fields of application. In this work, data collected from prediction markets for the FIFA World Cup 2006 is used to demonstrate their predictive power in the field of sports forecasting. For the first time, the prediction accuracy of play-money markets is compared to predictions based on historic soccer data as well as betting odds from professional bookmakers.

(II) How to design incentive schemes for play-money prediction markets?

Prediction markets can be used to provide incentives for information revelation. In real-money prediction markets you have to "put your money where your mouth is" (Hanson, 1990a). However, real-money prediction markets are illegal or at least highly regulated in most countries. Moreover, potential traders might be unwilling to invest their own money in prediction markets. Well-designed incentive schemes are thus needed to encourage participation and information revelation in play-money prediction markets. In this work, three different incentive schemes are compared with regard to their impact on the accuracy of predictions in a field experiment. In order to do so, predictions from three groups of traders with different and widely-used incentive schemes are compared with regard to their prediction accuracy in a field experiment. Subjects of the first group were paid a fixed amount, subjects of the second group were paid according to their ordinal rank within the group, and in case of the third group the subjects' payments linearly depended on their deposit value in the prediction market.

(III) How do traders' biases impact their trading behavior?

Prediction markets aggregate and reveal the information traders have. Individuals, however, exhibit substantial information processing or judgment biases. Markets which require probabilistic calculations and forecasts of future outcomes are particularly challenging with regard to the traders' information processing capabilities. Traders' biases may thus also affect their trading behavior in prediction markets and in doing so influence predictions based on trading prices. In financial markets, biases such as the home bias where investors allocate only a small fraction of their portfolio to foreign investments are a well-known phenomenon (e.g. French and Poterba, 1991). Traders in political stock markets are also buying and selling in a manner which is correlated with their party identification (Forsythe et al., 1992). This work studies how the traders' nationality impacts their holdings and their trading behavior in a FIFA World Cup 2006 prediction market.

1.3. Overview and Structure

The work at hand is structured into seven chapters. After the present introduction to this work, Chapter 2 gives a definition of prediction markets and explains their operational principle as well as their theoretical foundations. It also briefly discusses the key design elements of prediction markets which have to be considered by market engineers. Moreover, Chapter 2 presents current fields of application of prediction markets.

Chapter 3 describes a 2006 FIFA World Cup prediction market called STOCCER. Most of the data which is used to answer the research questions raised in the previous section comes from the STOCCER market. For this reason the FIFA World Cup 2006 itself, the contracts that were traded, the trading mechanisms, the incentive schemes, the group of traders, as well as the software platform are described in detail.

Chapter 4 examines the accuracy of prediction markets in general and in the field of sports forecasting in particular, more precisely for predicting the outcomes of soccer matches during the FIFA World Cup 2006. It thus answers the first research question raised in the previous section. The results show that play-money prediction markets outperform a random predictor and forecasts that are based on historic data about the success of national soccer teams. Moreover, prediction markets are on a level with betting odds from professional bookmakers which are known to be very accurate. Beyond the comparison of prediction accuracy, Chapter 4 also studies whether pure arbitrage opportunities existed in these markets and whether traders try to exploit illiquidity by taking on the role of market makers in prediction markets.

Afterwards, Chapter 5 studies the impact of different incentive schemes on prediction accuracy and thereby answers the second research question raised in the previous section. It elaborates on the question whether or not prediction markets with performance-related incentives perform better than markets with flat payments and how these performance-related incentives should be designed. This is of special interest when traders need to get paid for taking part in a prediction market, e.g. in the case of an internal market for company-specific predictions. The results show that the highest correlation between the relative frequency of outcome and trading prices is found in case of a rank-order tournament where traders are paid depending on their ordinal rank in a group of traders. Thus, tournaments with a handful of big winners winning big prizes work well. Somewhat surprisingly, the rank-order tournament even seems to beat the incentive scheme where the traders' payments are based linearly on their return in the market.

Chapter 6 analyzes how the traders' nationality impacts their holdings and their trading behavior in a FIFA World Cup 2006 prediction market. In doing so it answers the third research question from the previous section. Firstly, the chapter examines whether there is a correlation between the traders' nationality and the number of contracts they hold of different national teams. The results suggest that such a correlation does indeed exist. Secondly, it shows that traders tend to buy more contracts of their home country than traders from other countries do. In spite of these results predictions from these markets were surprisingly accurate.

Chapter 7 summarizes the contributions and discusses implications of this work. Finally, it proposes promising future fields of application for prediction markets and sketches future research questions that are closely related to those addressed in the work at hand.

1.4. Related Publications

Parts of this work have already been published and presented at research conferences. Concerning the results presented in Chapter 4, the comparison of play-money prediction markets to a random predictor and forecasts based on the historic data about the success of national soccer teams has already been published in Luckner et al. (2007).

Parts of Chapter 5 which examines the impact of different incentive schemes on prediction accuracy have already been published in Luckner and Weinhardt (2007). Furthermore, drafts and ideas of this research have been presented at various conferences: Dagstuhl Seminar "Negotiation and Market Engineering" 2006, Schloss Dagstuhl, Germany (Luckner, 2006b); 2nd Workshop on Prediction Markets, San Diego, USA (Luckner, 2007a), Doctoral Consortium of the 8. Internationale Tagung Wirtschaftsinformatik 2007, Karlsruhe, Germany (Luckner, 2007b); European Regional Meeting of the Economic Science Association (ESA), Nottingham, UK (Luckner, 2006a).

First ideas of the work on traders' biases discussed in Chapter 6 have been presented at the Group Decision and Negotiation Conference 2007, Mt. Tremblant, Canada

(Luckner, 2007c) and the 2007 Growth of Gambling and Prediction Markets Conference, Palm Desert, USA.

2. Prediction Markets

This chapter provides an overview of prediction markets. First, Section 2.1 explains what prediction markets are and how they work. Furthermore, the theoretical foundations of prediction market are outlined. Section 2.2 describes their key design elements before Section 2.3 gives an overview of several fields of application that have been reported in literature. Finally, Section 2.4 briefly summarizes the chapter.

2.1. Fundamentals of Prediction Markets

Throughout history business people have always tried to forecast the future to improve the performance of their companies. Commodity futures can be traced back to the Middle Ages when farmers and merchants faced the risk of price changes as a result of weather conditions or wars. In recent years, a relatively new approach for information aggregation has gained importance in the area of forecasting, namely prediction markets. Prediction markets bring a group of participants together and let them trade contracts whose payoff depends on the outcome of uncertain future events. The contracts thus represent a bet on the outcome of those future events. Once the outcome is known traders receive a cash payment in exchange for the contracts they hold.

Several studies describe how such markets have been applied for predicting future events or developments in the field of politics (Forsythe et al., 1992), sports (Luckner et al., 2007), medicine (Polgreen et al., 2007), or entertainment (Pennock et al., 2000). Moreover, companies like Siemens or Hewlett-Packard have employed prediction markets in order to improve their decision making (Chen and Plott, 2002, Ortner, 1997). This section contains a definition of what prediction markets are (2.1.1), a description of the operational principle of prediction markets (2.1.2) as well as the theoretical foundations of prediction markets (2.1.3).

2.1.1. Definition of Prediction Markets

In the academic literature, there is no universal definition of the term “prediction market”. Alternative terms used for the same concept include information markets, decision markets, idea futures, forecasting markets, artificial markets, electronic markets, and virtual stock markets. Figure 1 shows the number of research papers for different terms that are used to denominate the concept of prediction markets. The definition of prediction markets used in this work is based on Berg et al. (Berg and Rietz, 2003, Berg et al., 2003). According to this definition, prediction markets are defined as markets that are run for “the primary purpose of aggregating information so that market prices forecast future events” (Berg and Rietz, 2003, p. 3). Moreover, prediction markets can also serve as decision support systems by providing information about the current situation or by evaluating effects of decisions over time (Berg and Rietz, 2003, Hanson, 1999).

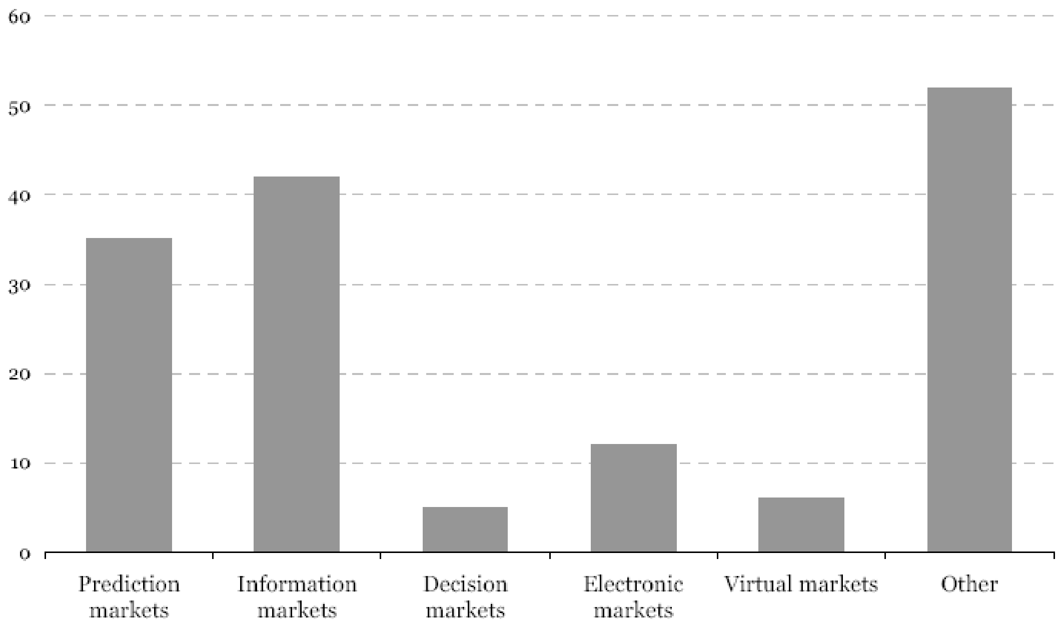

Figure 1: Number of research papers per term (Tziralis and Tatsiopoulos, 2007a)

Although prediction markets that are designed for information aggregation and revelation are at the focus of this work, the distinction between these markets and stock markets or betting markets can become fuzzy. In contrast to prediction markets, however, stock markets are established with the primary purpose of allocating resources, trading risk, and raising capital. Information aggregation is only a pleasant byproduct of stock markets while prediction markets are usually not substantial enough in size to allow for a considerable extent of risk sharing even though they may take on this role as interest and depth increase (Wolfers and Zitzewitz, 2004). Whereas contracts in stock markets are based on an underlying real asset, prediction markets create contracts which are linked to the outcomes of events but do not have any value by themselves. Betting markets, on the other hand, are first and foremost set up for entertainment and tend to trade risk that is intrinsically enjoyable. Thus, the primary purpose of a market can probably be seen as the main distinctive feature between prediction markets, betting markets, and stock markets.

2.1.2. Operational Principle of Prediction Markets

Prediction markets are a new form of financial markets where contracts whose payoff depends on uncertain future events are traded. Traders buy and sell contracts based on their expectations regarding the likelihood of future events. Trading prices thus reflect the traders' aggregated expectations on the outcome of uncertain future events and can be used to predict the likelihood of these events. The basic idea is that according to the efficient market hypothesis (Fama, 1970b) trading prices reflect all available information and the price mechanism serves as a means of aggregating the traders' collective expectations.

An example for the operational principle of prediction markets is shown in Figure 2. Suppose that the board of directors of a small deluxe car manufacturer needs reliable sales forecasts to adapt operational processes and minimize operational costs. All employees who have access to relevant information are given an initial endowment and access to the prediction market. Several contracts can be traded on this market. For example, the contract "500-600 cars in 2008" pays off 100 € if the company actually sells 500 to 600 cars in 2008; otherwise the pay-off is 0 €.

Assume that at a certain point in time the contract trades at a price of 45 €. In this case the trading price denotes that the probability that the car manufacturer will sell 500 to 600 cars in 2008 is assumed to be 45%. If a trader believes that the likelihood of selling 500 to 600 cars in 2008 is 70%, he should buy (sell) contracts for any price lower (higher) than 70 €. Thus, the trader would buy contracts at a price of 45 €.

As can be seen in this example a trader's dissent from the aggregated expectation would provoke a transaction and consequently usually change the trading prices. The trading mechanism automatically executes matching orders, i.e. buy and sell orders that are overlapping or placed at the same price. It is natural to assume that the higher a trader considers the probability of an event, the higher is both his reluctance to sell and his willingness to pay. Hence, the trading price gives some indication of how likely the traders as a group consider the event to occur. In this way, the trading price of the contract "500-600 cars in 2008" should reflect all the traders' information and can thus be interpreted as the probability of selling 500 to 600 cars in 2008.

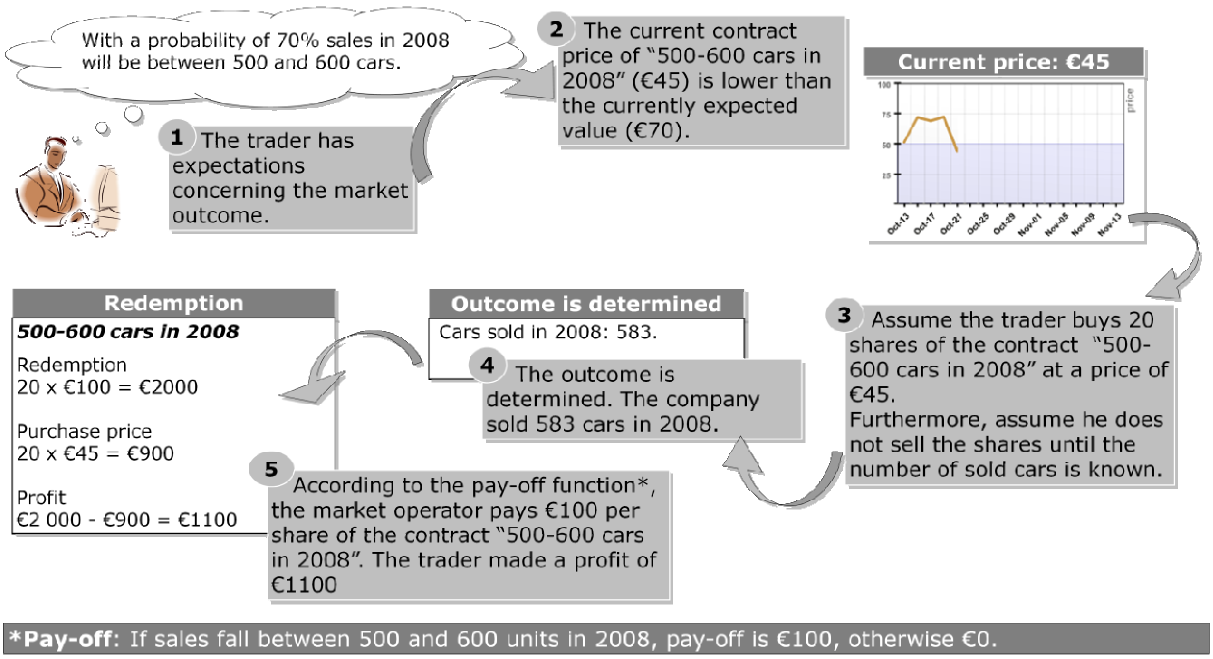

Figure 2: Operational principle of prediction markets

Depending on their performance, traders can either win or lose money. In the above-mentioned example, the trader bought 20 contracts "500-600 cars in 2008" at a price of 45 € and finally received a payment of 100 € per contract since the company indeed sold between 500 and 600 cars in 2008. Therefore, prediction markets motivate participation and well-designed incentive schemes motivate traders to reveal their beliefs instead of their preferences. To give an example, even an enthusiastic supporter of a deluxe car among the employees of the above-mentioned car manufacturer would rather not try to boost the sales forecasts of his favorite car since he would lose money in case he was overestimating sales figures.

2.1.3. Hayek and Efficient Market Hypotheses

The idea that trading mechanisms could be used to aggregate information dispersed among traders traces back to Hayek (Hayek, 1945). Hayek argued that planners in centrally-planned economies do not have enough information to calculate an optimal solution for resource allocation since central planners need information about all available resources and the preferences of people. He claimed that an efficient distribution of resources can only be maintained through the use of price signals in open markets. Accordingly, Hayek hypothesized that markets are the most efficient instrument to aggregate all the dispersed information of traders. Prices thus help to coordinate the separate actions of people.

“While the exact method by which information gets into the market is unknown” (Plott, 2000, p. 8), both theoretical and empirical research have found evidence that this process takes place. The efficient market hypothesis formulated by Eugene Fama states that stock “prices at any time ‘fully reflect’ all available information” (Fama, 1970b, p. 383). This implies that no additionally available information can be combined with efficient prices to improve the prediction accuracy of a market. Moreover, in financial markets it is impossible to consistently outperform the market by using any information that the market already knows. There are three common forms of market efficiency (Jensen, 1978). While the weak form efficient market hypothesis asserts that prices reflect all information contained in historic prices of the market, the semi-strong form efficient market hypothesis asserts that prices reflect all publicly available information. Of course, this also includes the past history of prices. Finally, the strong form efficient market hypothesis suggests that all relevant information known to anyone is reflected by the prices. The semi-strong form of the efficient market hypothesis is the accepted paradigm whereas there is evidence inconsistent with the strong form (Jensen, 1978).

Much of the enthusiasm for prediction markets derives from the efficient markets hypothesis due to the fact that contract prices reflect all information on the corresponding future event in an efficient prediction market and thus are the best predictor of future events. Information aggregation occurs when people can infer something from observing other traders’ believes and add that information to their own prior beliefs until there is a common knowledge equilibrium (McKelvey and Page, 1990).

Experimental research has tested the information aggregating properties of markets (e.g. Plott, 2000, Plott and Sunder, 1982, Plott and Sunder, 1988). In an experiment subjects traded contracts which paid 200 if the state was Y and 400 if the state was X with probabilities of 0.75 and 0.25. During so called informed states, some insiders knew the state of the world. As can be seen in Figure 3 prices in these markets converged to the correct value when insiders were present and for the most part to the expected value of 250 if none of the traders were insiders. Thus, these markets were able to collect and broadcast information held by some of the traders (Plott, 2000).

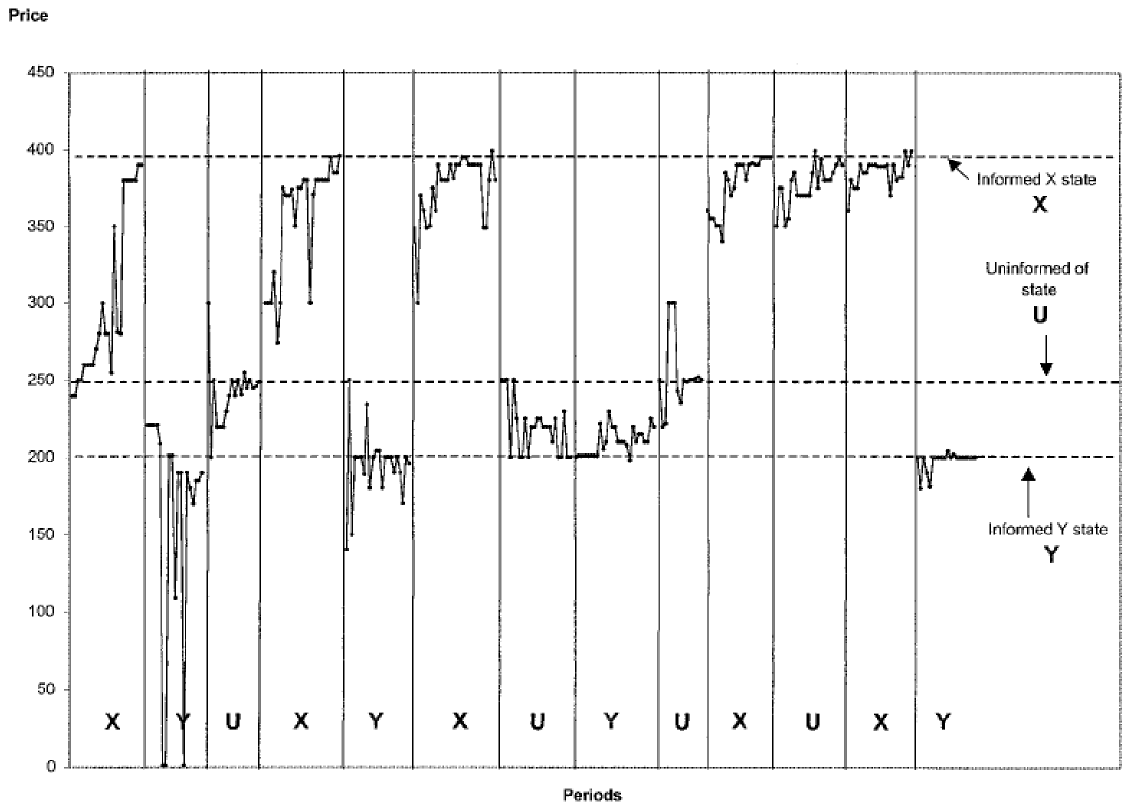

Figure 3: Information of insiders is revealed in trading prices (Plott, 2000)

In real-world scenarios, however, knowledge is usually dispersed among traders. Consequently, the question arises whether markets can aggregate this dispersed information. Therefore, in another experiment every subject was given partial, private information. Collectively, the traders had almost perfect information regarding the correct state. The results show that information aggregation did also occur in this case (Plott, 2000).

2.2. Key Design Elements of Prediction Markets

Before studying more advanced applications of prediction markets, it is necessary to gain a basic understanding of their key design elements. Like any market, prediction markets have to be designed and implemented very carefully in order to ensure that they are suitable for aggregating traders' information (Weinhardt et al., 2003, Weinhardt et al., 2006a). The key design elements comprise the specification of *contracts* traded in a prediction market, the *trading mechanism*, and the *incentives* provided to ensure information revelation (Spann and Skiera, 2003). Moreover, diversity of information is required in order to provide a basis for trading (Wolfers and Zitzewitz, 2004). Disagreement among traders is desirable and the selection of *traders* is thus also considered a key design issue (Tziralis and Tatsiopoulos, 2007b). The following subsections describe these design elements in more detail.

2.2.1. Contracts

Prediction markets can be used to predict absolute numbers such as sales in a fiscal year, relative numbers such as market share, and the occurrence or non-occurrence of a particular event such as a natural disaster in a certain geographic region. The transformation of the forecasting goals into contracts should be carried out in a way that the contracts are clear and easily understood. Wolfers and Zitzewitz distinguish three basic types of contracts, namely "winner-take-all", "index", and "spread" contracts (Wolfers and Zitzewitz, 2004).

A winner-take-all contract pays of certain sum of money if an event occurs and doesn't pay anything otherwise. As a result, the price of a winner-take-all contract can be interpreted as the traders' aggregated expectation of the probability of the occurrence of a future event, for example the probability of a team winning a soccer match (Wolfers and Zitzewitz, 2006a).

Index contracts link the payoff directly to a number such as the percentage of the popular vote that a candidate will receive in a political election. Thus, the trading price for such a contract represents the mean value that the market assigns to an outcome.

Spread betting establishes a cutoff that defines the occurrence of an event such as whether a candidate receives more than a certain percentage of the popular vote. In consequence it reveals the market's median expectation if contracts are designed in such a way that winners double their money while losers do not receive any payment. This is only a fair bet in case the payoff is as likely to occur as not.

These are only the basic types of contracts and real-world prediction markets are making use of all kinds of variations of them. One important aspect with regard to the design of contracts is to provide contingency resolutions if the underlying facts that determine the contract change or if the results become non-verifiable. To give an example, a prediction market could be employed to predict product sales in an accounting year. What happens if the company decides to stop selling the product due to a liability suit before the end of the year? Suchlike cases have to be considered when setting up a market.

2.2.2. Trading Mechanisms

The most integral aspect of any trading platform is how buyers and sellers are matched. The most widely used trading mechanism in the field of prediction markets is the continuous double auction (CDA). Alternative mechanisms are call auctions (CA), dynamic pari-mutuel markets (DPM), as well as market scoring rules (MSR). These mechanisms are briefly described in the following. Table 1 summarizes their advantages and disadvantages with regard to three desirable properties of trading mechanisms for prediction markets, namely continuous incorporation of information, guaranteed liquidity, and avoidance of financial risk for the market operator (Pennock, 2004).

Table 1: **Comparison of trading mechanisms**

	CDA	**CA**	**DPM**	**MSR**
Continuous information incorporation	Yes	No	Yes	Yes
Guaranteed liquidity	No	No	Buying: yes Selling: no	Yes
Risk for operator	No	No	No	Yes (bounded)

Continuous Double Auction (CDA)

So far, the continuous double auction (CDA) is the most commonly used trading mechanism in prediction markets. In case of a CDA, as known e.g. from continuous trading at the electronic financial trading system Xetra of Deutsche Börse AG[8], traders submit buy and sell orders which are executed immediately if they are executable against orders on the other side of the order book (Madhavan, 1992); if not, orders are queued in an order book and remain there until they expire, are matched with a counteroffer, or are removed. Usually, orders are executed according to price/time priority, i.e. buy orders with a higher limit and vice versa sell orders with a lower limit take priority. In case several orders were placed with the same limit the orders which were submitted earlier are executed first.

One of the main advantages of using a CDA is the fact that markets with a CDA pose no financial risk for the market operators. Since it only matches willing traders all markets can be implemented as a zero-sum game (Spann and Skiera, 2003). As a consequence, this mechanism is especially popular among real-money exchanges. The Iowa Electronic Markets (IEM), for example, have started using the CDA in their markets in the late 1980ies (Forsythe et al., 1992). Moreover, the CDA allows for continuous information incorporation into prices and consequently traders are capable of quickly reacting to events in case of liquid markets.

However, with few traders the markets may suffer from illiquidity, e.g. when many shares are traded or few traders are active in the market. Offers can then not be matched with counteroffers and therefore the bid-ask spread can be huge or order queues are empty (Hanson, 2003). Since most prediction markets have fewer participants than traditional financial markets, this limitation is particularly relevant for them. The trading mechanisms that are discussed in the following draw on different approaches to address the thin market problem.

Call Auction (CA)

In financial markets call auctions are used as an alternative trading mechanism. The electronic trading system Xetra of Deutsche Börse AG, for instance, uses a hybrid system of continuous double auctions and call auctions. While orders are executed immediately in continuous markets they are accumulated for simultaneous execution at a pre-determined point in time according to a priority rule, e.g. the principle of the highest executable volume, in call auctions (for an overview, see (Madhavan, 1992)).

[8] http://www.xetra.com

Liquidity in illiquid low-volume markets can consequently be accumulated and focused on pre-determined execution times. Although trading in illiquid markets is also not possible as long as there is no matching counter offer, call auctions make it more difficult to move trading prices and thus influence the price formation with small transactions compared to continuous markets. Just like the CDA mechanism call auctions also pose no financial risk for the market operator. Due to periodic trading in call auctions, however, new information is not reflected immediately in trading prices.

STOCCER[9] was one of the very few prediction markets implementing call auctions. Thus, it remains an open question whether call auctions are suitable trading mechanism for prediction markets. Results from the STOCCER market suggest that traders prefer to trade in continuous markets. The trading activity measured by the number of trades per day was higher in case of the CDA than in the call auction market (Geyer-Schulz et al., 2007). This result is in line with the findings in financial markets where traders that are faced with the decision of choosing either form of market also prefer continuous markets (Kalay et al., 2002).

Dynamic Pari-Mutuel Market (DPM)

Dynamic pari-mutuel markets (DPM) are a hybrid between the above-mentioned continuous double auction and pari-mutuel markets which are e.g. traditionally employed for horse-race betting. In pari-mutuel markets money goes into a central pool and is later divided among the winners. This provides infinite liquidity and circumvents the thin-market problem of double auctions. There is no need for a matching offer from another trader. But one shortcoming of pari-mutuel markets is that there is no incentive to buy contracts early, especially not if new information is expected before the market closes. Purchasing contracts will also inform other traders. As a result, it is the best strategy to wait until the last possible moment to buy. “Prices” in consequence cannot be considered a reflection of current information.

Pennock has developed the DPM mechanism in order to combine the infinite liquidity of pari-mutuel markets with a trading mechanism in which prices continuously react to new information (Pennock, 2004). The DPM offers infinite buy-in liquidity and thus acts as a one-sided market maker always offering to sell at some price and moving the

[9] http://www.stoccer.com

price according to demand. Prices are computed using a price function which can differ depending on the properties that are desired. The DPM also does not exhibit any risk of losses for the market operator due to its redistribution of money. Moreover, it allows traders to lock in gains or limit losses by selling contracts in a CDA market. Selling still has to occur through a CDA mechanism because there is no market maker accepting sell offers. Nevertheless, traders can always "hedge-sell" by buying the opposite outcomes (Pennock, 2004).

The DPM has been implemented in Yahoo's Tech Buzz game[10], a prediction market for high-tech products, concepts, and trend (Mangold et al., 2005).

Market Scoring Rules (MSR)

Hanson's market scoring rule (MSR) acts like a two-sided market maker that also provides infinite liquidity for the sell side of the market with a variable but bounded maximum loss that can be regarded as a subsidy for the market (Hanson, 2003). Market scoring rules can be thought of as sequentially used proper scoring rules. An MSR maintains a probability distribution over all events. Any trader who believes the probabilities are wrong can change the current prediction by replacing it with a new prediction as long as the trader agrees to pay off the most recent person. If traders improve the prediction by moving the prices into the right direction they can expect a positive payoff, otherwise they will lose money. New information is hence reflected immediately.

This MSR has already been implemented by exchanges such as InklingMarkets[11] or the Washington Stock Exchange[12].

2.2.3. Incentives

Appropriate incentives schemes are required to motivate participation and to ensure information revelation in prediction markets. The traders' remuneration is crucial for the success of a market and consequently a key design element. Previous research in the field of prediction markets has shown that play-money as well as real-money markets can predict future events to a remarkable degree of accuracy. One relevant

[10] http://buzz.research.yahoo.com
[11] http://inklingmarkets.com
[12] http://www.thewsx.com

question is how much difference it actually makes whether markets are run with real money or with play money (Wolfers and Zitzewitz, 2004). Even though one might intuitively expect the performance of play-money markets to be worse than the performance of real-money markets, some have argued that “play money exchanges may even outperform real-money exchanges because ‘wealth’ can only be accumulated through a history of accurate predictions” (Wolfers and Zitzewitz, 2004). A study of the predictions of the 2003 NFL football season has shown that the real-money market TradeSports and the play-money market NewsFutures predicted outcomes equally well (Servan-Schreiber et al., 2004).

Due to the legal restrictions on gambling many prediction markets nowadays rely on play money. Some traders may be intrinsically motivated but even in play-money markets the market operators can provide incentives such as a flat fee for participation or prizes for the largest play-money fortunes to remunerate traders. Chapter 5 discusses selected incentive schemes for play-money markets and their impact on the accuracy of prediction in more detail.

2.2.4. Traders

In the end, prediction markets only work if traders with relevant information join the market and trade. Market operators in consequence have to make sure they select traders with relevant information. One straightforward approach could be to invite experts who have access to information concerning the under study claims. This was usually done in corporate prediction markets, e.g. by Hewlett-Packard and Siemens (Chen and Plott, 2002, Ortner, 1997). These markets had only between 20 and 60 traders and companies have repeatedly cited motivating employees to participate as an obstacle to a more wide-spread use of prediction markets (Wolfers and Zitzewitz, 2004). However, inviting experts only has at least two downsides.

Firstly, most prediction markets have very few participants compared to traditional financial markets. As a result, it is hard to fill an order book in a CDA market. The lack of offers to buy and sell limits the incentive for traders to reveal new information because they will have difficulty finding a trading partner. Replacing the widespread CDA by another trading mechanism is one approach to ensure that traders can profit from new information without having to find a trading partner. This downside can therefore be by-passed with a suitable market design.

Secondly and even more important, it is rather unlikely that there is a lot of disagreement among fully rational experts trading in a market. Disagreement about likely outcomes, however, is required to encourage trading (Wolfers and Zitzewitz, 2004). Overconfident traders as well as an increase in noise trading should actually improve the accuracy of trading prices because this increases the rewards to informed trading – provided informed traders have deep pockets relative to the volume of noise trading. This is consistent with earlier research on prediction markets demonstrating that markets aggregate information and produce efficient outcomes despite biased individual traders (Forsythe et al., 1999). Also, experimental results confirm that manipulators in prediction markets are unable to distort price accuracy (Hanson et al., 2006).

Instead of limiting the pool of traders to knowledgeable experts one should thus try to attract more traders. If traders self-select to join a market they usually have relevant information about and considerable interest in the under study claims. Nevertheless, one should avoid running markets on topics where insiders may possess substantially superior information or where information is concentrated on very few people. Such markets have historically attracted very little attention (Wolfers and Zitzewitz, 2004). Equilibrium prices may in this case not accurately reflect the true probabilities because informed traders do not completely reveal their information. This can be explained by the fact that few informed traders can frequently benefit from fluctuating trading prices repeatedly and thus do not reveal their information at once. The example shows that trading mechanisms such as the CDA may be ill-suited for small scale markets because the market design is not incentive compatible (Ledyard, 2006).

2.3. Fields of Application

This section gives an overview of previous fields of application of prediction markets that have been reported in the literature. Since it is all but impossible to consider the totality of earlier applications, the list of applications given in Table 2 was compiled based on an extended literature review which was recently published in the Journal of

Prediction Markets in an attempt to collect the totality of academic work related to prediction markets (Tziralis and Tatsiopoulos, 2007a).

Table 2: Fields of application of prediction markets

	Market	**Focus**	**Reference**
Political stock markets	Iowa Electronic Markets	US presidential elections, non-US elections (e.g. Austria, France, Korea, Germany)	Berg et al. (2001), Berg et al. (1996), Berg et al. (1997), Berg and Rietz (2003), Berg et al. Berg and Rietz (2006), Bondarenko and Bossaerts (2000), Erikson and Wlezien (2006), Forsythe et al. (1994), Forsythe et al. (1992), Forsythe et al. (1999), Fowler (2006), Kou and Sobel (2004), Oliven and Rietz (2004)
	UBC election stock market	Provincial and federal elections in Canada	Antweiler and Ross (1998), Forsythe et al. (1995), Forsythe et al. (1998)
	Swedish EU PSM	Swedish 1994 EU referendum	Bohm and Sonnegard (1999)
	GEM 90, GEM 91, GEM 94, GEM 98	Federal and regional elections in Germany	Brüggelambert (2004)
	Wahlstreet, Wahlboerse	State elections in Germany	Hansen et al. (2004)
	Passauer Wahlbörse	Federal elections in Germany	Beckmann and Werding (1996)
	The Political Stock Market	Federal and provincial elections in Germany	Franke et al. (2006), Franke et al. (2005)
	NP02, TE03	National assembly and regional elections in Austria	Huber and Hauser (2005)
	"Die Presse" Election Market	Elections for the national assembly in	Filzmaier et al. (2003)

	Market	Focus	Reference
		Austria 2002	
	Austrian Political Stock Market	Austria's membership in the EU, federal elections, governing coalition	Ortner et al. (1995)
	PAM94	European Parliament and municipal councils in the Netherlands	Jacobsen et al. (2000)
Sports prediction markets	TradeSports	Worldwide sports prediction market, e.g. baseball, soccer, football	Chen et al. (2005), Rosenbloom and Notz (2006) , Servan-Schreiber et al. (2004)
	NewsFutures	Sports (e.g. baseball, football, soccer), political elections	Chen et al. (2005) , Rosenbloom and Notz (2006), Servan-Schreiber et al. (2004)
	World Sports Exchange	Football, baseball, hockey, basketball etc.	Debnath et al. (2003)
	Betfair	Soccer, tennis, horse racing, etc.	Smith et al. (2006)
Other applications	Hollywood Stock Exchange	Box office performance of movies	Gruca et al. (2003), Pennock et al. (2001b), Pennock et al. (2001a)
	CMXX	Success of movies, music CD's and video games in Germany	Skiera and Spann (2004)
	Economic Derivatives	Retail sales, GDP, international trade balance, growth in payrolls	Gürkaynak and Wolfers (2006)
	Tech Buzz Game	High-tech products, concepts, and trends	Mangold et al. (2005)
	Foresight Exchange	Future developments in science and technology	Pennock et al. (2001b), Pennock et al. (2001a)

Table 2 comprises all applications of the prediction market concept that were reported in journal articles, books or book chapters, and conference proceedings papers referenced in the aforementioned literature review. Pure lab experiments where signals are e.g. drawn from an urn were not taken into consideration. The applications were grouped into three categories: political stock markets, sports prediction markets, and other applications. Due to the fact that most of the longest running prediction markets were originally set up to forecast political elections or the outcome of sports tournaments, academic research has largely concentrated on political stock markets and sports prediction markets. The following subsections provide some more information on the three categories of applications.

2.3.1. Political Stock Markets

Beside early introductory articles by Hanson (Hanson, 1990a, Hanson, 1990b, Hanson, 1992), most of the literature on prediction markets up until 1998 is on political stock markets. The most cited and earliest application of a political stock market on the internet, the Iowa Electronic Markets (IEM[13]), was established in 1988 by the University of Iowa. The IEM were designed to give students a hands-on experience in trading and to study market dynamics. The first academic article on the IEM was published in 1992 (Forsythe et al., 1992). IEM focussed on US presidential and state elections, but the platform was also used to run political stock markets on elections e.g. in Austria, France, Korea, and Germany. Predictions derived from IEM trading prices have been more accurate than their natural benchmark, namely polls, although traders exhibit biases (Berg et al., 2001, Forsythe et al., 1999). Moreover, trading prices react extremely quickly to new information (Berg and Rietz, 2006). In the meanwhile the IEM are not only used for predicting the outcome of political elections but also in order to predict e.g. economic indicators. Beside predicting uncertain future events the IEM were also studied as a decision support system where decisions are made based on trading prices (Berg and Rietz, 2003).

Other political stock markets in Canada (e.g. Antweiler and Ross, 1998), Sweden (Bohm and Sonnegard, 1999), Germany (e.g. Beckmann and Werding, 1996), and Austria (e.g. Ortner et al., 1995) have been set up with a similar research focus.

[13] http://www.biz.uiowa.edu/iem/

Furthermore, these markets were also used to study manipulation in prediction markets (Hansen et al., 2004). All in all, political stock markets have in many cases outperformed traditional polls (Berg et al., 2001). Due to this reason they have received quite a lot of attention in the media and several publishing houses have already been running their own markets (Filzmaier et al., 2003).

2.3.2. Sports Prediction Markets

Sports prediction markets like Betfair.com [14], the World Sports Exchange [15], NewsFutures [16], and TradeSports [17] are among the most popular prediction markets. These markets focus on forecasting the outcome of sports tournaments and events. Among popular sports are e.g. baseball, soccer, football, hockey, basketball, tennis, and horse racing. Although NewsFutures, for instance, does also operate markets on politics, financial markets, or the movie business, contracts on sports events are usually the most popular topics. Earlier studies on sports prediction markets show that these markets provide at least as accurate predictions as experts (Chen et al., 2005, Servan-Schreiber et al., 2004). In accordance with the efficient market hypothesis game events are quickly resulting in changes of trading prices. Smith et al. (2006) find that markets on UK horse racing exhibit both weak and strong form of market efficiency.

One precondition for exploiting the potential of prediction markets is to provide incentives for participation and information revelation. Therefore, prediction markets such as the IEM require real-money investment from traders. In case of the IEM these investments are limited to a maximum amount of US$ 500. As was already mentioned in Section 2.2.3 two articles in the field of sports prediction markets, however, show that there is no significant difference in terms of prediction accuracy between play-money and real-money prediction markets (Rosenbloom and Notz, 2006, Servan-Schreiber et al., 2004).

2.3.3. Other Applications

Nowadays, prediction markets are increasingly employed in innovative fields of application beyond political stock markets and sports prediction markets. One popular

---

[14] http://www.betfair.com
[15] http://de.wsex.com
[16] http://us.newsfutures.com
[17] http://www.tradesports.com

example is the Hollywood Stock Exchange (HSX[18]), a prediction market where traders forecast box office revenues of films, both for opening weekends and beyond. CMXX.com was a similar market operated in Germany to predict the success of movies, music CD's, and video games (Skiera and Spann, 2004). Pennock et al. (2001a) demonstrate that trading prices in the HSX movie markets are good predictors of the box office performance of movies. Based on these forecasts the movie industry can then make decisions on how to allocate advertising based on expected box office revenues. This demonstrates how companies can use prediction markets to make better informed decisions.

Apart from predicting box office revenues markets can be used broadly for predicting the success of all kinds of new products (Gruca et al., 2003). Successful examples for such markets are the simExchange[19], a market for predicting the sales of console hardware and upcoming video games, or an internal market run by Eli Lilly to find out which drugs will be most successful (Kiviat, 2004). Prediction markets can thus be seen as an alternative to traditional marketing research techniques. Spann et al. (2007) show that prediction markets are also useful for identifying lead users with superior abilities to forecast the market success of new products. Their idea is that lead users perform better than average traders on prediction markets. The percentage of lead users among the best performing traders is similar to the percentage found in survey-based screening.

Another interesting field of application is the prediction of macroeconomic data such as retail sales, GDP, international trade balance, and the growth in payrolls. For this purpose a market called "Economic Derivatives[20]" was launched in 2002. A first analysis shows that the expectations reflected in trading prices are similar to survey-based predictions (Gürkaynak and Wolfers, 2006).

Up to now, prediction markets were mostly applied to forecast events in the near future. Determining the payoff of a particular contract is then straightforward as soon as the outcome becomes known. Yet, some of the earlier research also proposes the use of prediction markets for forecasting events in the distant future (Hanson, 1992). One

[18] http://www.hsx.com
[19] http://www.thesimexchange.com
[20] http://www.economicderivatives.com

market for predicting long-term developments in science, technology, and other fields of public interest is the Foresight Exchange[21] (Pennock et al., 2001a). Contracts traded in this market range from technical to socio-political issues. Another market for long-term predictions which is exploiting the potential of prediction markets to continuously update trading prices is the Tech Buzz Game[22]. Yahoo Research sponsors this market which lets traders predict the technologies that internet users will be searching the web for in the future (Mangold et al., 2005). One market could be trading contracts on rival technologies such as web browsers. These contracts then pay a weekly dividend relative to the number of search requests. In the long term, the market closes if the topic becomes uninteresting and the contracts will then be liquidated for cash. One of the goals of the Tech Buzz Game is to test dynamic pari-mutuel markets which were discussed in Section 2.2.2 in the field.

Other prominent examples of companies using prediction markets internally are Hewlett-Packard where traders produced more accurate forecasts of printer sales than the company's forecasting team (Chen and Plott, 2002) or Siemens where software developers predicted the completion date of a huge software project (Ortner, 1997).

2.4. Summary

This chapter gave a short introduction to the field of prediction markets. The term prediction market as it is understood in this work was defined. Moreover, the theoretical foundations as well as the operational principle of prediction markets were described.

Like any other market, prediction markets have to be designed carefully. The key design elements, namely the contracts traded in a market, the trading mechanisms, incentives for traders to participate and reveal their expectations, as well as the selection of traders have been introduced. In addition, several design alternatives for each of these design elements have been briefly discussed.

At the end of the chapter, the main fields of application of prediction markets that have been reported in literature were presented. So far, academic literature for a large part

[21] http://www.ideosphere.com
[22] http://buzz.research.yahoo.com

focussed on political stock markets although numerous companies have already made use of internal corporate prediction markets. This can probably be explained by the fact that companies do oftentimes not want to make their experiences public. Concerning the field of sports prediction markets there are up to now only very few research papers.

3. STOCCER – A 2006 FIFA World Cup Prediction Market

This chapter describes a 2006 FIFA World Cup prediction market called STOCCER. Most of the data which is used to answer the research questions in the following three chapters comes from the STOCCER prediction market. Section 3.1 describes the FIFA World Cup 2006 itself before Section 3.2 presents the STOCCER exchange including its key design elements as well as information about traders and the trading activity. Section 3.3 outlines the trading software which was used as the basis of the STOCCER prediction market. Finally, Section 3.4 briefly summarizes the chapter.

3.1. The FIFA World Cup 2006

The most important soccer tournament worldwide in 2006, the FIFA World Cup, was held in Germany from June 9th to July 9th 2006 with 32 participating national teams which had qualified for the tournament. The tournament was organized in two stages – a group stage and a knock-out stage. All in all, 48 matches were played in the group stage and 16 in the knock-out stage, resulting in a total of 64 matches.

In the group stage the teams played round robin in eight groups of four to qualify for the knock-out stage. The winning team of a match received three points, the losing team received zero points, and in case of a draw after 90 minutes each team received one point. The two most successful teams in each group advanced to the knock-out stage. If two or more teams achieved the same number of points the direct comparison, i.e. the results of the match(es) against each other, was used as a tie-breaker. Further subordinate tie-breakers are the difference between the numbers of goals scored and received, the total number of goals scored in the group stage, the FIFA country coefficient from the FIFA world ranking, and finally tossing a coin.

In the knock-out stage, which started on June 24, the winning team of a group played the second of one of the remaining groups. All the matches in the knock-out stage were played in a sudden death system. Additionally, one game was played for the third place between the losers of the two semi-final games. In case of a draw after regular time in the knock-out stage the match was continued for an extra time of two times fifteen minutes. If a match was still not decided after extra time, there were penalty shootouts. The winner of a match in the knock-out stage advanced to the next round. Figure 4 shows all the 16 matches from the knock-out stage of the FIFA World Cup 2006.

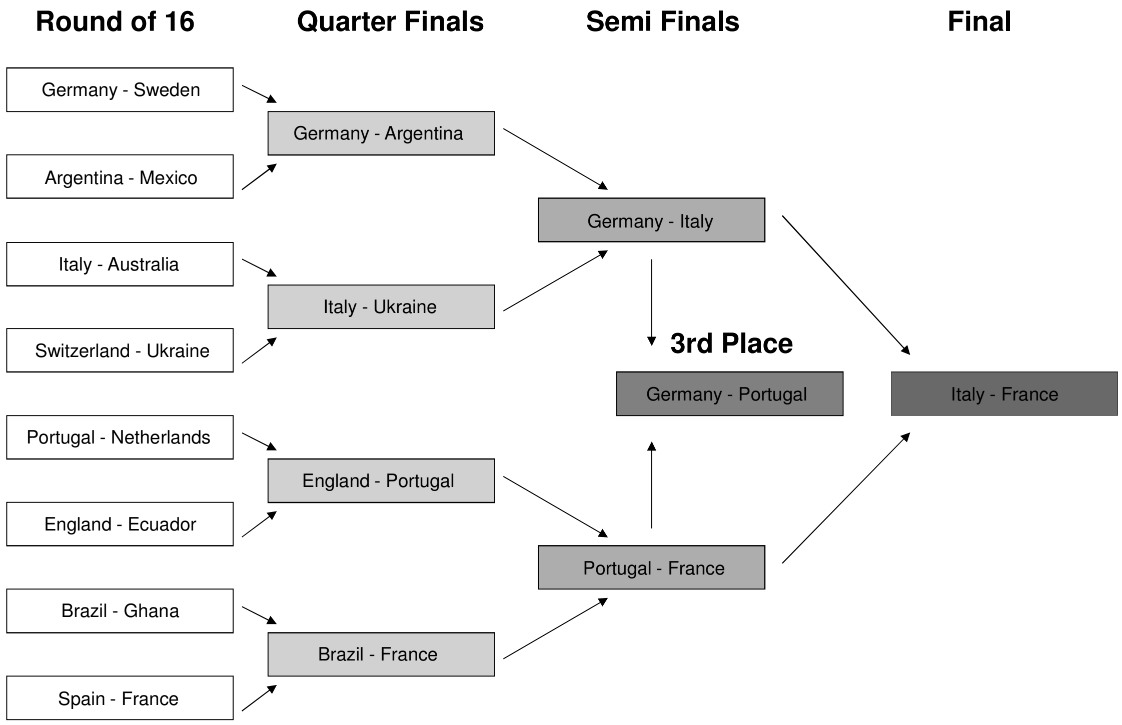

Figure 4: Knock-out stage of the FIFA World Cup 2006

The tournament was won by Italy, defeating France in a penalty shootout after extra time finished in a draw. Germany defeated Portugal to finish third. After the sometimes surprising 2002 tournament, the FIFA World Cup 2006 was dominated by traditional soccer powers. Six former champions took part in the quarter finals with Ukraine and Portugal remaining as the only relative outsiders.

3.2. The STOCCER Exchange

STOCCER was operated before and during the 2006 FIFA World Cup in order to predict the outcome of the tournament, the outcome of particularly exiting matches, and the tournament's top goal scorer. In total, more than 1.700 traders registered with the play-money prediction market STOCCER[23]. The first market started on May 15th 2006 and ran until the end of the FIFA World Cup on July 9th 2006. The trading platform

[23] http://www.stoccer.com. The STOCCER project was funded by the German Federal Ministry for Education and Research under grant number 01HQ0522.

was open to the public 24 hours a day, 7 days a week. On average, there were more than 1,600 trades per day with a total number of about 90,000 trades. The continuous increase in the number of registered users as well as the development of the trading activity through time is illustrated in Figure 5. The upsurge in the number of users and the number of trades per day around June 9th 2006 can without much doubt be explained as follows. First of all, the opening match took place that day and consequently there was a lot of interest in the tournament. Furthermore, several newspaper articles on the STOCCER exchange were published at that time and the markets were thus made known to a larger audience.

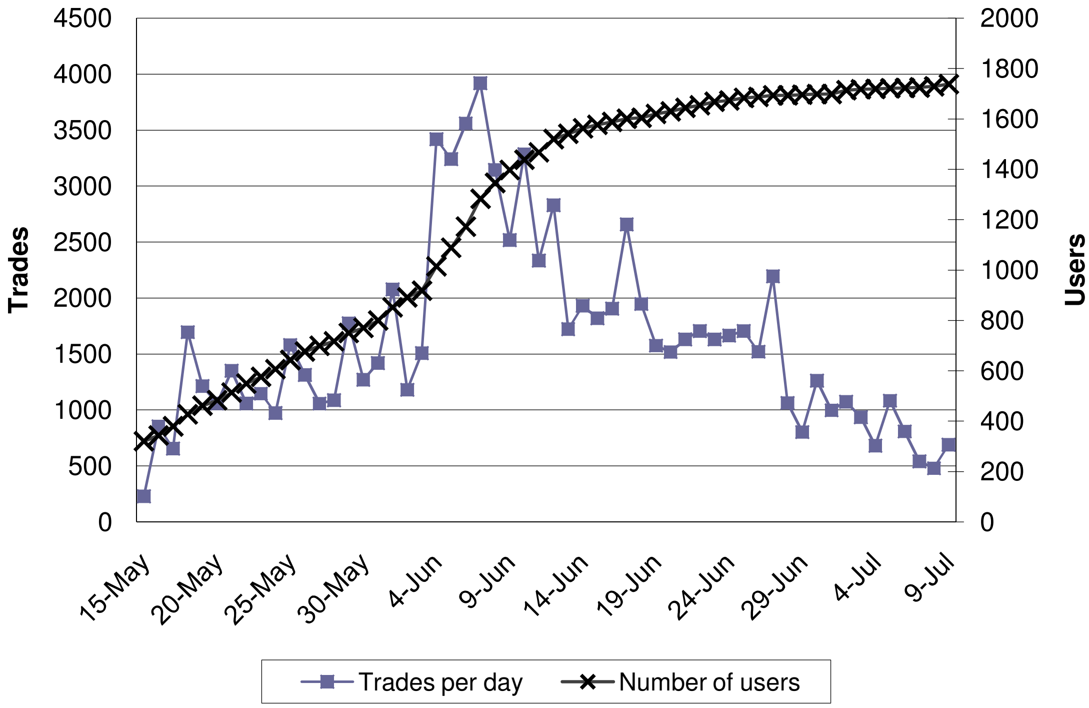

Figure 5: Number of users and trading activity over time

The following subsections describe the key design elements of our markets, i.e. the contracts that were traded, the trading mechanisms, the incentive schemes, and the group of traders, in more detail.

3.2.1. Contracts

In total, we ran 19 markets – 16 markets for the 16 matches in the final rounds starting with the round of sixteen, two markets to predict the tournament's top goal scorer, and the so called championship market where shares of all the 32 national teams taking part in the FIFA World Cup 2006 were traded. These three types of markets are also shown in Table 3 with some more information on the number of contracts that were traded in each of the markets, market start and end time, as well as information on how the contracts were valued at the close of the market.

Table 3: Markets operated during the FIFA World Cup 2006

Type	Number of contracts	Payoff	Start time	End time
Championship	1 per country (32)	World champion: 50 Vice-WC: 30 Semi finals: 20 Quarter finals: 10 Round of 16: 5 Otherwise: 0	May 15th 2006	July 9th 2006
Match	3 per match: team A wins, team B wins, tie after 2nd half	Event occurred: 10 Otherwise: 0	2 days before the matches	At the end of the matches
Goal scorer	Fluctuating	Top goal scorer: 100 Otherwise: 0	June 6th 2006	July 9th 2006

In case of the first type of markets, namely the championship market, the 32 contracts of the national soccer teams were valued as follows at the close of the market: 50 virtual currency units for the world champion, 30 for the runner-up, 20 for all the teams dropping out in the semi finals, 10 for those dropping out in the quarter finals, and 5 for all those dropping out in the round of 16. All shares of the remaining 16 teams were worthless in the end. The championship market started about three weeks before the first match of the FIFA World Cup 2006 and was closed immediately after the final on July 9th 2006. It was the only market which was online for the complete time period of the world championship.

More than 1,260 traders submitted orders to this market and in total there were more than 80,000 trades. The total number of trades per contract is depicted in Figure 6. Among the most heavily traded contracts are mainly traditional soccer powers such as France, Germany, Brazil, and Argentina. One reason for the relatively high number of trades in case of “Angola” could be that contracts in the order input mask were sorted alphabetically and the contract of Angola was thus listed first.

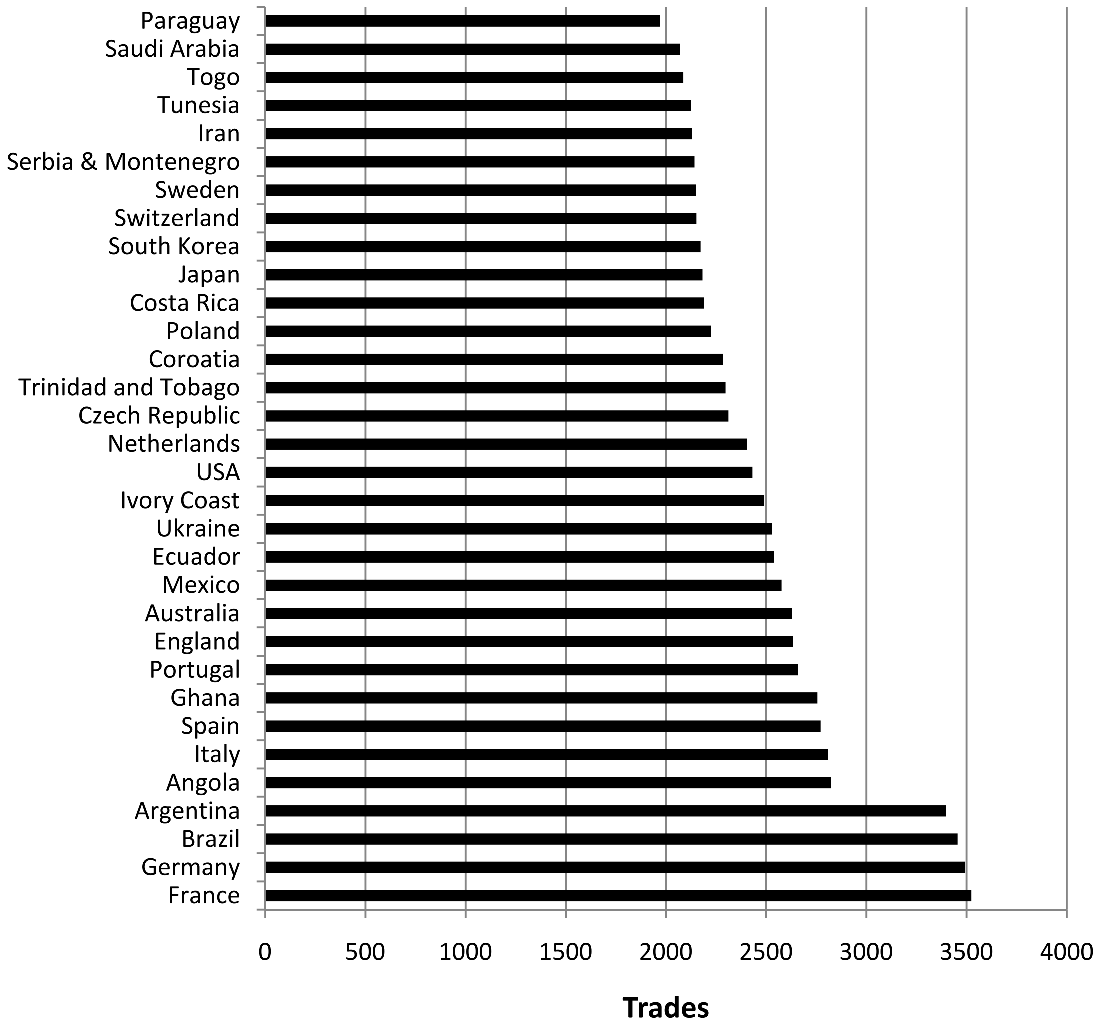

Figure 6: Number of trades in the championship market

The second type of markets, namely the match markets, focused on predicting the outcome of matches in the final rounds. For the 16 matches in the final rounds there were three contracts per match. This is because the following three possible outcomes for every match were defined: Either one of the two national teams won or there was a draw after the second half. The third contract (“draw”) was introduced although there were no draws possible in the final rounds of the tournament. The reason for this was that overtimes and penalty shootouts were not considered as their outcomes can be regarded as more or less unpredictable. This is also rather common in case of sports betting with professional bookmakers. Trading started two days before the matches and was stopped immediately after the second half of the matches. The contract corresponding to the event that actually occurred was valued at 10 virtual currency units after the match; the other two contracts were worthless.

Data on the trading activity in the 16 match markets is given in Figure 7 which shows the number of traders as well as the number of trades per match market. On average, there were about 110 traders per market who submitted orders during the two days the markets were open. With 120 trades only "Switzerland-Ukraine" was the match with the smallest number of trades. The most liquid market was the semi final "Portugal-France" with nearly 900 trades. On average, there were about 450 trades per match market.

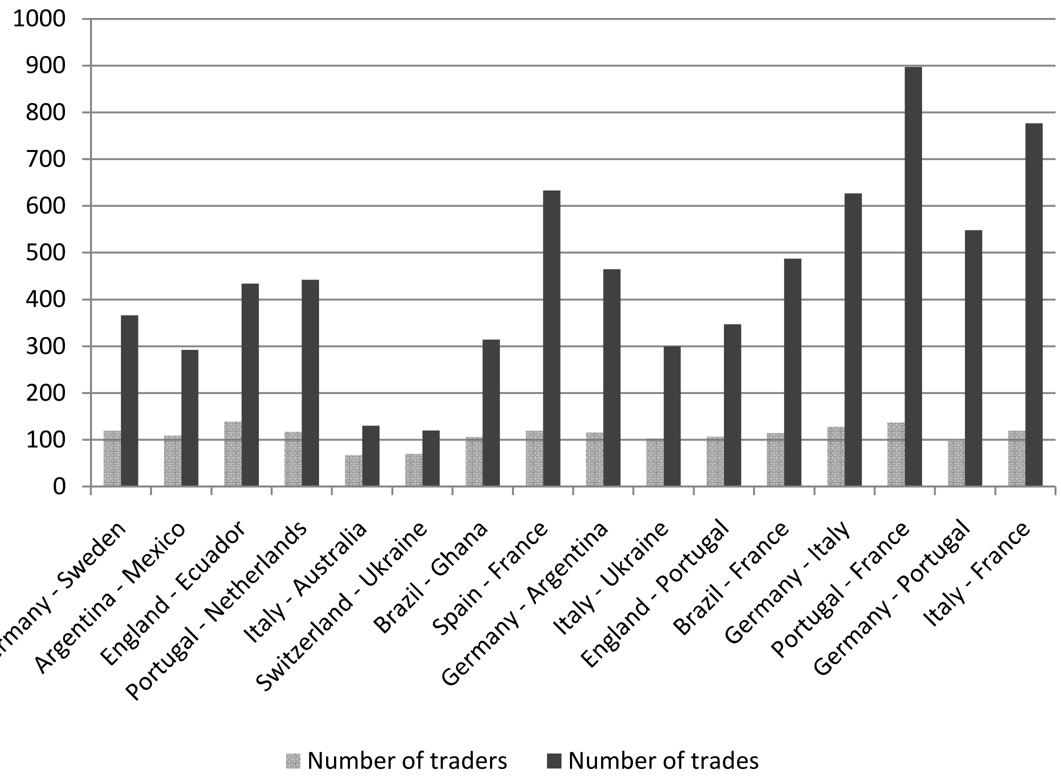

Figure 7: Trading activity in the match markets

The idea behind the third type of markets, namely the two goal scorer markets, was to predict the top goal scorer of the whole tournament. The contract of the top goal scorer was valued at 100 virtual currency units; all other contracts were valued at 0. If there were several top players with the same number of goals, these would have been valued at 100 virtual currency units divided by the number of those players. Initially, the goal scorer market was started with a pre-determined set of players on June 6th 2006. Additionally, there was a contract "other", which was split into two contracts as soon as a player which had so far not been traded in the market scored his third goal. In this case, a contract corresponding to the new player was introduced to the market. If a trader had shares in "other" in his deposit at this point in time, he received an additional contract of the new player automatically.

In order to study the impact of the trading mechanism on the prediction accuracy and the trading behavior there were two goal scorer markets – one market with a continuous double auction and a second market with a call auction. Traders were free to choose any of the two markets for buying and selling their contracts in individual players. Figure 8 depicts the number of trades over time in both markets.

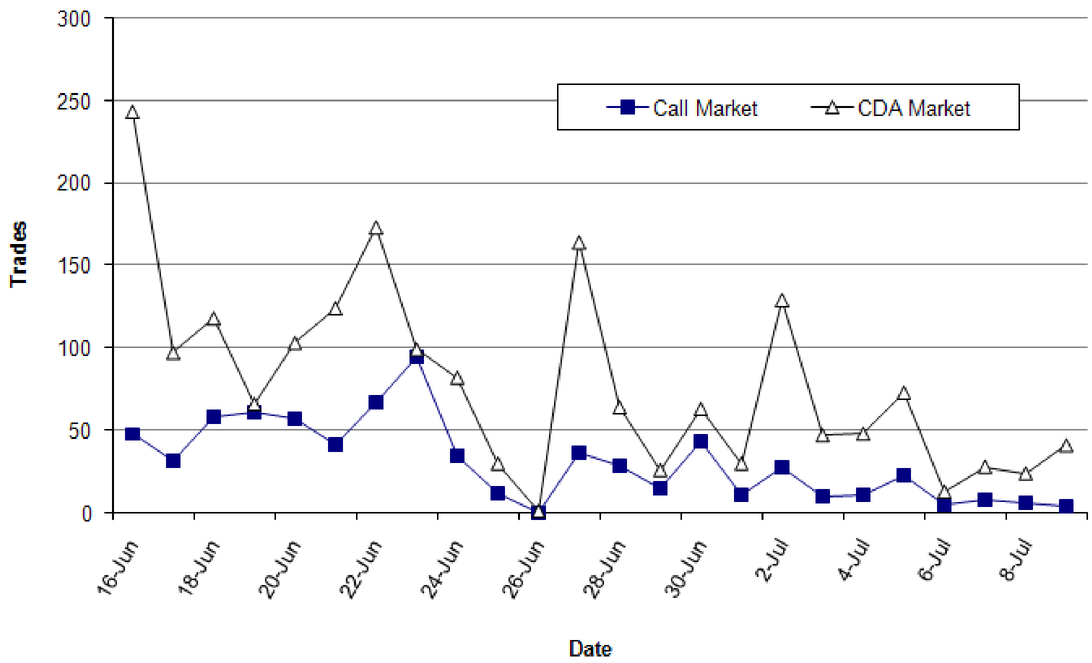

Figure 8: Distribution of trades per day over time

It is obvious that the trading activity measured by the number of trades per day was higher in case of the CDA market than in the call market. On average, there were more than 78 trades per day in the CDA compared to 31 trades per day in the call auction. In total, there were 1886 trades in the CDA market compared to 738 trades in the call market. For some reason traders seem to prefer trading in the CDA market. Looking at the number of traders that had at least one trade in the respective market the CDA market with 197 traders also outnumbers the call market with 179 traders.

3.2.2. Trading Mechanisms

Concerning the financial market design, two different trading mechanisms were used in STOCCER – continuous double auctions (CDA) and a call auction. These two trading mechanisms were already roughly explained in Section 2.2.2. The only non-CDA market was one of the two goal scorer markets. Since this market is of no particular importance for answering the research questions addressed in this work it is not described in more detail. All of the other markets, i.e. the championship market, the 16 match markets, as well as the second goal scorer market, employed a CDA in combination with limit orders.

Upon registration each trader was assigned 100 shares of each contract traded in any of the markets as well as a cash account of 100,000 virtual currency units and was thus able to trade instantly. Additional shares were issued by means of so called basic portfolios (Forsythe et al., 1992). A basic portfolio contains one share of every contract which is traded in the respective market. The portfolio price equals the sum of the payoffs for one share of every contract in a market and was e.g. 10 virtual currency units in case of the match markets. It thus corresponded to the payoff for correctly predicting the outcome of a match. Buying and selling portfolios from and to the market operators was therefore risk free for traders and possible at any time while the markets were operating.

Traders submitted offers to buy (bids) or offers to sell (asks). Bids and asks were maintained in queues with a price/time priority, i.e. they were first ordered by price and then by time. Offers remained in the queues until (i) they were withdrawn by the traders, (ii) their lifetime as defined by the trader had expired, or (iii) they were matched with a counter offer. The trades were automatically executed as soon as bid and ask prices in the respective queues were overlapping. When a bid was submitted at a price equal to or exceeding the current minimum price in the ask queue, a trade was executed at the ask price. Analogously, when a sell offer was submitted at a price equal to or less than the current maximum price in the bid queue, a trade was executed at the bid price. In case there were two or more offers at the same price, the earliest offer submitted to the market was executed first. Since the system did not analyze the traders' identities a trader could also trade against himself. Short sales were disallowed by the system. Moreover, submitting offers with insufficient funds in the cash account as well as offers to sell when the trader's portfolio did not contain the corresponding number of shares in a contract were prevented.

3.2.3. Incentives

In contrast to traditional betting exchanges for sports events the prediction market STOCCER was operated as a play-money market. Setting up a real-money sports prediction markets is currently not legal in Germany. Instead of investing real money every trader had an initial endowment of 100,000 virtual currency units as well as 100 shares of each contract. The only extrinsic incentives for traders to join the market and reveal their expectations were a ranking of their user names on the STOCCER web page and a lottery of prizes. The overall TOP-100 traders, i.e. the 100 traders with the highest deposit value after the final of the FIFA World Cup on July 9th 2006, took part in a final lottery where the first prizes were shares of the "Garantiefonds UniGarant Deutschland (2012)" investment fund with a value of 3,000, 2,000, and 1,000 Euro. Traders thus had a rather strong incentive to be among the 100 traders with the highest deposit value. In addition, we weekly raffled an iPod among the 20 most active traders of the preceding week.

The most successful trader was able to increase his deposit value by almost 900% between May 15th 2006 and July 9th 2006. At the other extreme, several traders lost almost 100% of their initial deposit value. General terms and conditions were used to prevent traders from creating multiple user accounts and trading against themselves in order to transfer cash from one account to another. Traders were not allowed to register more than once. Furthermore, the use of any kind of software for automated actions was prohibited. Several traders violated these terms and conditions and were disqualified.

3.2.4. Traders

Participation in STOCCER was voluntary. In total, more than 1,700 traders enrolled in the prediction market. During the registration process traders provided information about their gender, age, and country of origin. Traders were predominantly male and quite young compared to the total population of their countries of origin. Almost 89% of the traders were male. Table 4 shows the traders' age distribution. Traders of age 30 and younger account for almost 57% of the total number of traders.

Table 4: Age distribution of traders

Age	Number of traders	Proportion of traders	Year of birth
<= 20	96	5.26%	>= 1987
20-25	486	26.64%	1982-1986
26-30	454	24.89%	1977-1981
31-35	232	12.72%	1972-1976
36-40	155	8.50%	1967-1971
41-45	137	7.51%	1962-1966
46-50	111	6.09%	1957-1961
51-55	69	3.78%	1952-1956
51-60	38	2.08%	1947-1951
>= 60	46	2.52%	<= 1946

Since STOCCER was operated and made known in Germany traders coming from this country also formed the largest group of traders. Overall, traders originated from 72 different countries around the world. As can be seen in Figure 9 about two thirds of the traders were German.

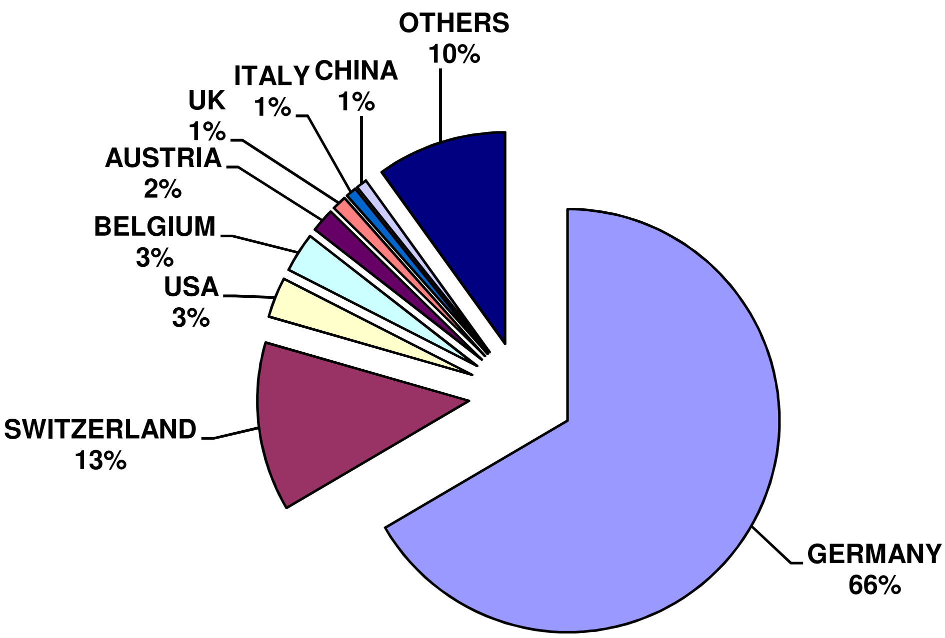

Figure 9: Traders' country of origin

Other countries with a substantial number of traders were Switzerland (235 traders), USA (56 traders), Belgium (55 traders), Austria (33 traders), UK (20 traders), China (15 traders), and Italy (15 traders).

After the FIFA World Cup all of the traders were asked to complete a brief web-based survey to provide descriptive information amongst others about their knowledge and interest in soccer as well as their experience in securities trading. 74 traders completed this survey. Three quarters of these traders saw 16 or more matches during the FIFA World Cup live on TV. 13 out of the 74 traders saw even more than 45 matches on TV during a period of four weeks only. Thus, they seem to be rather enthusiastic about soccer. Several traders also appear to be rather experienced in securities trading. More than 55% of the traders who completed the survey hold a portfolio of securities and about 10% of them trade quite a lot in financial markets, i.e. they conduct more than 20 transactions per year. 27% of the traders completing the survey were even familiar with the concept of prediction markets and had already participated in other prediction markets.

3.3. The Trading Software

In addition to the key design elements of the STOCCER prediction market described in the previous section one also has to design the web-based trading software as well as the facilities provided for obtaining information about the traders' accounts, the different markets, offers, and trades from a technical point of view. STOCCER had to meet numerous functional and non-functional requirements such as running several prediction markets simultaneously, each of them in multiple languages, or enabling different trading mechanisms for different markets. A fairly flexible platform was needed since it should be easy to reuse in other fields of application such as e.g. market research. Due to the large number of users the software platform also had to be scalable.

In order to fulfill all the requirements the STOCCER trading software was based on two existing trading platforms and thus integrated the functionality of these systems. The two platforms were the political stock market PSM[24], a field-tested platform which

[24] http://psm.em.uni-karlsruhe.de

was in the past primarily used for predicting the outcomes of political elections (Franke et al., 2005), and meet2trade[25], a generic electronic trading platform that realizes innovative trading features such as bundle trading and enables traders to individually configure their own electronic market (Weinhardt et al., 2006b, Weinhardt et al., 2005). The most liquid market, i.e. the championship market, was operated based on the PSM while all the match markets and the goal scorer markets were run with the meet2trade trading platform. Depending on the market a user wanted to trade in he was forwarded to a trading screen provided by either of the two trading platforms.

The traders of course should not take notice of the fact that STOCCER was built on two existing platforms. Thus, a web interface with exactly the same look and feel for both trading platforms was implemented. An example of the main trading screen is shown in Figure 10.

Market information available to traders included the accumulated bids at the highest three bid prices, the accumulated asks at the lowest three ask prices, the last trading price, and charts showing the price history of all contracts. Moreover, a short description of the market comprising the respective payoff function was shown as part of the trading screen. An alert service informed traders via e-mail in case individual price limits which had been predefined by the respective trader were exceeded. Available account information for individual traders included the number of shares held in each contract, the balance of the cash account, the total value of their deposit, a list of outstanding buy and sell orders, as well as a list of trades.

A ranking of all the traders sorted by their deposit value, i.e. the balance of their cash account plus the value of the contracts they held at the specific point in time, was not part of the trading screen but was separately displayed on the STOCCER web portal www.stoccer.com. This portal also provided more information on the prizes traders could win, the operational principle of the prediction market including a tutorial and frequently asked questions, as well as up-to-date soccer news related to the FIFA World Cup 2006. All the information from the trading screen and the portal was available in four languages, namely German, English, French, and Spanish.

[25] http://www.meet2trade.com

Figure 10: Trading screen of STOCCER

Because the PSM and meet2trade are not based on the same technology, the two trading platforms were integrated on the database level. As can be seen in Figure 11 both systems accessed the same PostgreSQL database. All the required data such as user data was shared by the PSM and meet2trade, so that a trader had to register only once and was then granted access to both of the underlying trading platforms. The dividing rule between the two platforms was the type of contract which was traded. This means that contracts traded in the championship market – which was operated based on the PSM – were not at the same time traded in other markets run by meet2trade and vice versa. Nevertheless, the traders' deposits had to be integrated because both platforms made use of the same cash account. Coordinating the trading activity was consequently required in the sense that e.g. the total volume of a trader's buy orders in both systems was not allowed to exceed the amount of money in his cash account. Both trading platforms also provided market administration tools, e.g. for adding new markets and contracts.

As Figure 11 shows the common PostgreSQL database[26] was operating on one physical machine and was accessed from the two machines which were used to run the two trading platforms PSM and meet2trade (m2t). The STOCCER web portal was built up using the TYPO3 Content Management System[27] and ran on a fourth machine. A separate MySQL database[28] was used to store the content of the portal.

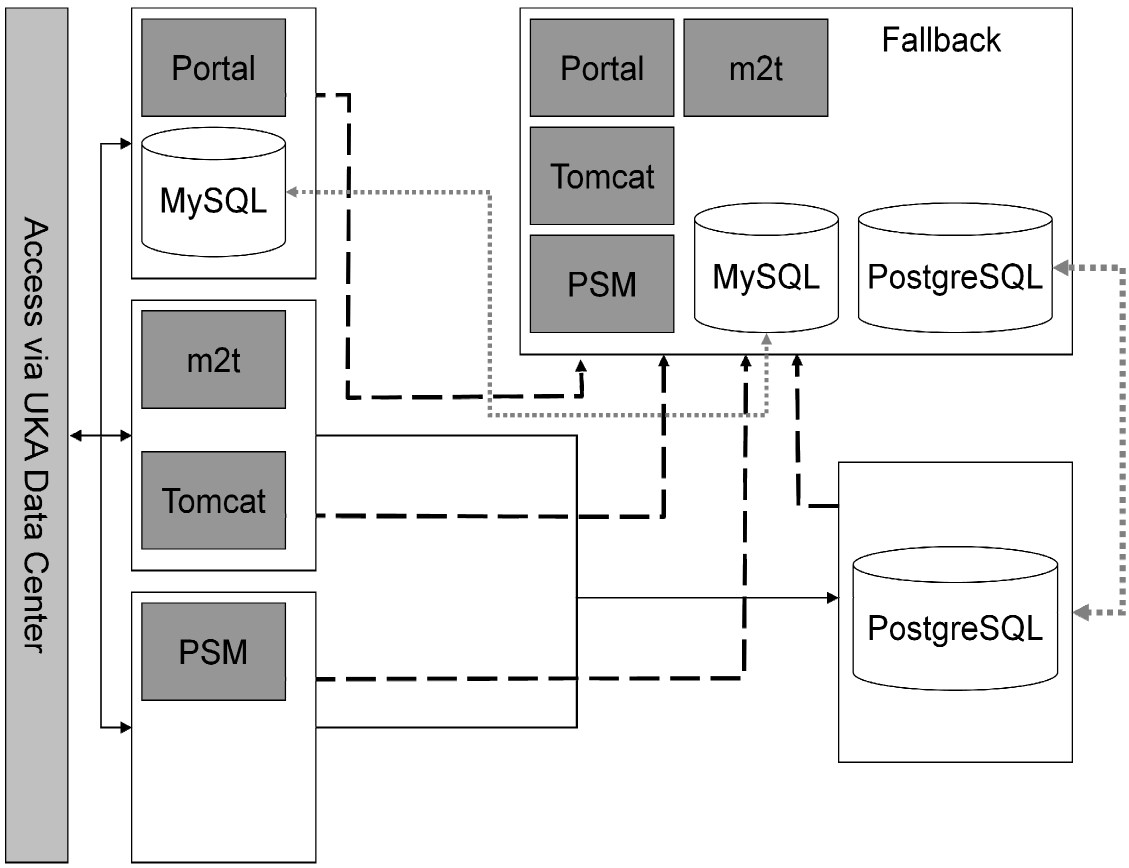

Figure 11: Hardware and software architecture of STOCCER

Running these software systems on four different machines was required to cope with the system load. In order to guarantee the continuous operational availability of the STOCCER trading software a fifth machine was ready to take over the tasks performed by one of the forth other machines at any time. For this purpose the data from the two

[26] http://www.postgresql.org/
[27] http://typo3.org/
[28] http://www.mysql.com/

databases had to be replicated on the fifth machine because the data might otherwise be lost forever or at least be temporarily unavailable.

3.4. Summary

This chapter described the sports prediction market STOCCER which was operated by the Universität Karlsruhe (TH) in collaboration with the University of Frankfurt before and during the FIFA World Cup 2006 in Germany. Most of the data which is used to answer the research questions in the following three chapters comes from the STOCCER markets.

STOCCER was presented along the lines of the key design elements that have already been introduced in Chapter 2. Markets were run in order to predict the outcome of the tournament as a whole, the outcome of specific matches, and the tournament's top goal scorer. All but one market were using a continuous double auction trading mechanism. In general, one can presumably say that the markets did not suffer from illiquidity compared to other prediction markets run by academic institutions. In order to provide incentives for traders from all over the world – with the largest share coming from Germany – prizes were raffled among the most successful and the most active traders.

The chapter concluded with a brief description of the trading software which was used for running the markets and which has in the meanwhile also been used in other fields of application.

4. Empirical Evidence of Prediction Accuracy

Accurate forecasts are essential in various fields of application such as supply chain management, technology forecasting, economic forecasting, or product forecasting. Product forecasting which aims at predicting the level of success of a new product is particularly suitable to demonstrate the importance of accurate predictions for corporations. Predictions in this context must take into account aspects like product awareness, distribution channels, prices, competitive alternatives, and consumer behavior. Inaccurate predictions can result in considerable costs for a company and may weaken its position compared to its competitors in the market. In particular, predicting demands for products with short product life cycles or predicting sales in unstable market situations is challenging (Spann and Skiera, 2003). In the past, flop rates of new products were high, oftentimes surpassing 50% (e.g. Urban and Hauser, 1993). Thus, reducing flop rates by means of more accurate predictions can have a huge impact on the profits and increase the competitive advantages of a company.

The application of prediction markets is a new approach which can be used for product forecasting and beyond. One of the main reasons for the emergence of prediction markets is that markets have done well in comparison with other forecasting methods (Hanson, 2006). Horse race markets, for instance, beat horse race experts (Figlewski, 1979) and Oscar markets beat columnists (Pennock et al., 2001b). Usually, prediction markets tend to perform at least as well as the single best individual, without requiring a priori knowledge of whom that individual is (Surowiecki, 2004). Prediction markets are thus considered to provide a method to improve prediction accuracy compared to traditional forecasting methods.

This chapter provides evidence of their prediction accuracy in general and in the field of sports forecasting in particular. Earlier empirical research substantiates the predictive power of markets relative to traditional forecasting methods such as expert opinions or polls in various fields of application. Data collected from the play-money prediction market STOCCER for the FIFA World Cup 2006[29] is used to empirically compare the prediction accuracy of sports prediction markets to (i) random predictors, (ii)

[29] STOCCER was presented in detail in the chapter 3.

predictions that are based on historic soccer data about the success of national soccer teams, as well as (iii) betting odds from professional bookmakers. Thus, this chapter contributes to the literature by providing the first empirical comparison of play-money prediction markets with predictions based on historic data or betting odds in the field of sports forecasting.

The idea behind using these three benchmarks is the following: Forecasts of prediction markets are driven by the traders' information and expectations. These forecasts are worthless if they do not result in better predictions than randomly drawing possible outcomes. Thus, random predictors are used as a first benchmark to evaluate the prediction accuracy of the STOCCER markets. Beside historic data, traders also consider current information available to them as well as ongoing developments within the course of the tournament. Using predictions based on the historic success of national soccer teams as a second benchmark allows for examining whether markets are superior to these predictions by incorporating additional information. Within the scope of this research, the FIFA world ranking[30] is used as it is calculated based on pure historic data. Betting odds serve as a third benchmark since they are well-established in sports and known for being very efficient (cp. Gandar et al., 1998, Pope and Peel, 1989). Fixed-odds betting differs from prediction markets since the odds are determined by experts, i.e. the bookmakers, and bettors can only decide whether or not to place a bet at the given price. In prediction markets, in contrast, prices reflect the traders' aggregated expectations and can be changed by any trader with deviating expectations.

Prediction markets should work well if they are efficient, and in efficient markets, one does not expect arbitrage opportunities to be persistent. Beyond the comparison of prediction accuracy, this chapter therefore also studies whether pure arbitrage opportunities existed in STOCCER. Moreover, market liquidity can become an issue in prediction markets since new information is potentially not immediately reflected in trading prices and traders might also lose interest in the markets if those are illiquid. It is therefore analyzed whether traders try to exploit illiquidity by taking on the role of market makers in prediction markets.

[30] http://www.fifa.com/worldfootball/ranking/

The remainder of the chapter is structured as follows. Section 4.1 presents related work on the analysis of the prediction accuracy of markets in general as well as markets for sports forecasting in particular. Section 4.2 then describes how predictions for the outcome of specific soccer matches are derived from trading prices in the STOCCER markets. It also presents the data which is used to compare play-money prediction markets to alternative forecasting methods. In Section 4.3, the prediction accuracy of the STOCCER prediction markets is analyzed by comparing the predictions to a random predictor, predictions derived from the FIFA world ranking, and betting odds. Furthermore, it is analyzed whether pure arbitrage opportunities existed and whether traders acted as market makers. Section 4.4 discusses the results before Section 4.5 briefly summarizes the main findings of this chapter.

4.1. Related Work

A large body of earlier research in the field of prediction markets focuses on evaluating their prediction accuracy in absolute terms or relative to alternative forecasting methods. The results reported in Section 4.1.1 are at large convincing and provide evidence that non-sports prediction markets do well in comparison with other forecasting methods. Since the focus of this chapter is on the prediction accuracy of sports prediction markets, previous studies from this field of application are discussed separately in Section 4.1.2.

4.1.1. Non-Sports Prediction Markets

As already mentioned in Section 2.3.1, a large share of the literature on prediction markets treats political stock markets. It is thus not surprising that several articles on the prediction accuracy of political stock markets in absolute terms and relative terms compared to opinion polls have been published since 1988. In 1988, prediction markets were for the first time introduced as an alternative to traditional opinion polls in order to predict the outcome of the presidential election in the US. The 1988 US-presidential market almost perfectly predicted the candidates' vote share (Forsythe et al., 1992). The actual vote share and the vote share predicted via the market are depicted in Table 5[31]. Bush's vote share was predicted accurately to a tenth. Moreover, the market underestimated the Dukakis vote share by two tenths of a percentage point and

[31] The last column contains the predicted shares of the vote as predicted by the IPSM market.

overestimated the combined strength of all remaining candidates by six tenths of a percentage point. Altogether, this prediction was almost perfect.

Table 5: 1988 US presidential elections and forecasts (Forsythe et al., 1992)

Candidate	Total votes	Vote share	IPSM forecast vote share
Bush	48,138,478	53.2	53.2
Dukakis	41,114,068	45.4	45.2
Third party	1,219,240	1.4	2.0

Since then, over the last four presidential elections prior to 2004 the Iowa Electronic Markets (IEM) have predicted vote shares with an average absolute error of around 1.5 percentage points while the final Gallup poll erred by 2.1 percentage points (Wolfers and Zitzewitz, 2004). In addition, using data from four US presidential elections, Wolfers and Zitzewitz (2004) showed how prediction accuracy improves over time. The horizontal axis of Figure 12 shows the number of days until the election and the vertical axis quantifies the average absolute forecast error, i.e. the average absolute deviation between predicted and actual vote share. Accordingly, prediction accuracy improves as information is revealed and reflected in trading prices prior to the election.

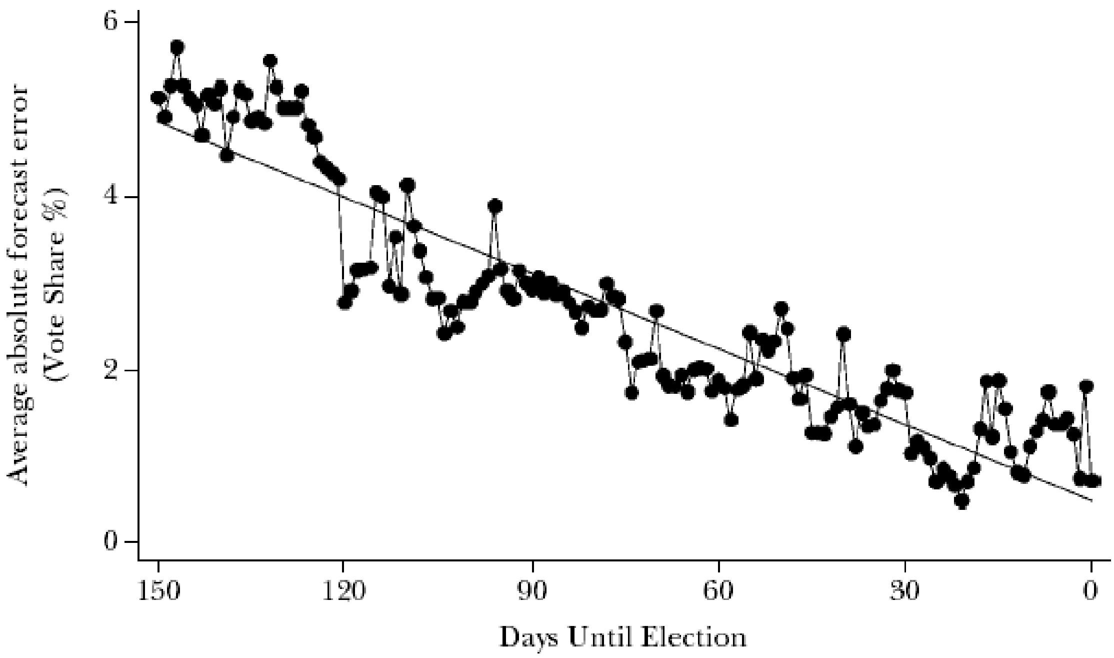

Figure 12: Information revelation through time (Wolfers and Zitzewitz, 2004)

Concerning the prediction accuracy of political stock markets in absolute terms, Berg et al. (2001) compared the predicted to the actual outcomes for vote-share and seat-shares markets. Figure 13 plots the predicted versus the actual election outcomes for 49 markets run in 13 countries. As can be seen, most of the elections are close to the 45-degree line which represents perfect accuracy. Predictions are consequently often very close to the actual outcome.

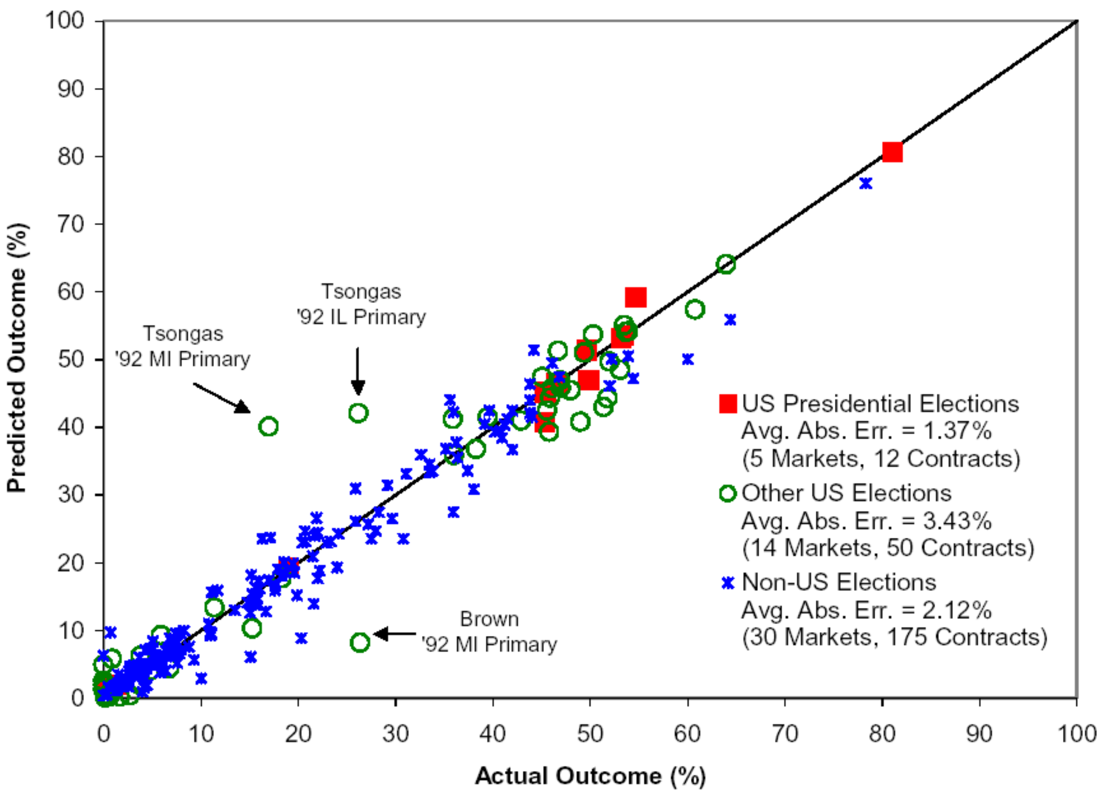

Figure 13: Predicted vs. actual outcomes in political markets (Berg et al., 2001)

However, political stock markets are not only accurate in absolute terms but also relative to election polls. They beat election polls in many cases. Figure 14 compares the mean absolute error of political stock markets and election polls for a total of 15 political elections in the United States and several European countries. Errors of polls in this case are average errors across major polls from the last week before the election. Berg et al. (2001) found that predictions of markets are closer to the actual outcome than polls in 9 out of 15 cases. The average poll error is 1.93 percentage points across all elections while the average error of the market is only about 1.5 percentage points. In a more extensive study, Berg et al. (2003) compared the accuracy of the Iowa

Electronic Markets IEM to traditional polls by analyzing 596 national polls between 1988 and 2000. The survey reveals that the prices of the IEM outperform polls in 76 percent of the time.

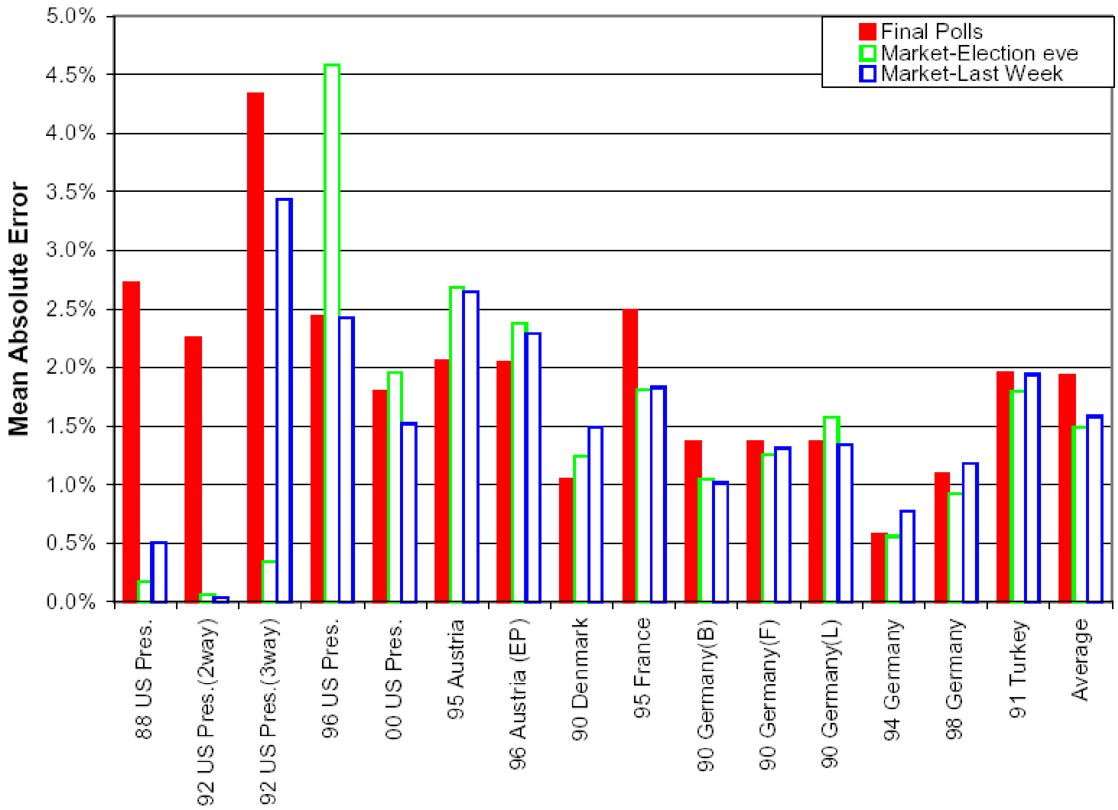

Figure 14: Political stock markets compared to polls (Berg et al., 2001)

Other studies present similar findings from other domains beside political stock markets, demonstrating that prediction markets perform well compared to traditional forecasting methods like surveys, opinion pools, or expert judgments. In case of an internal market at Hewlett-Packard (HP), for example, Chen and Plott (2002) found out that the prediction market beat the official sales forecasts of the company in 6 out of 8 cases for which official forecasts were available. Hence, the markets performed better than traditional methods employed inside HP. It is noteworthy to mention that, in contrast to the IEM, only a small number of people were selected for participating in the HP markets. Additionally, markets were operated over short periods of time only. Prediction markets consequently seem to work even under these circumstances.

Another well-known prediction market which was already mentioned in Section 2.3.3 is the Hollywood Stock Exchange (HSX). The HSX allows traders to trade on, for instance, the opening weekend performance and total box office returns of movies. Figure 15 shows that the predictions of the opening weekend box offfice success of the HSX have been remarkably accurate (Wolfers and Zitzewitz, 2004) in the past. Pennock et al. (2001a) also find that the HSX forecasts are a good predictor for the opening weekend and the four week box office returns. In 2007, the HSX correctly predicted seven out of eight Oscar winners in the top categories and thus seems to work almost perfectly (Lamare, 2007). For the 2000 Oscars, the HSX has beaten the individual and average forecasts of five experts (Pennock et al., 2001a).

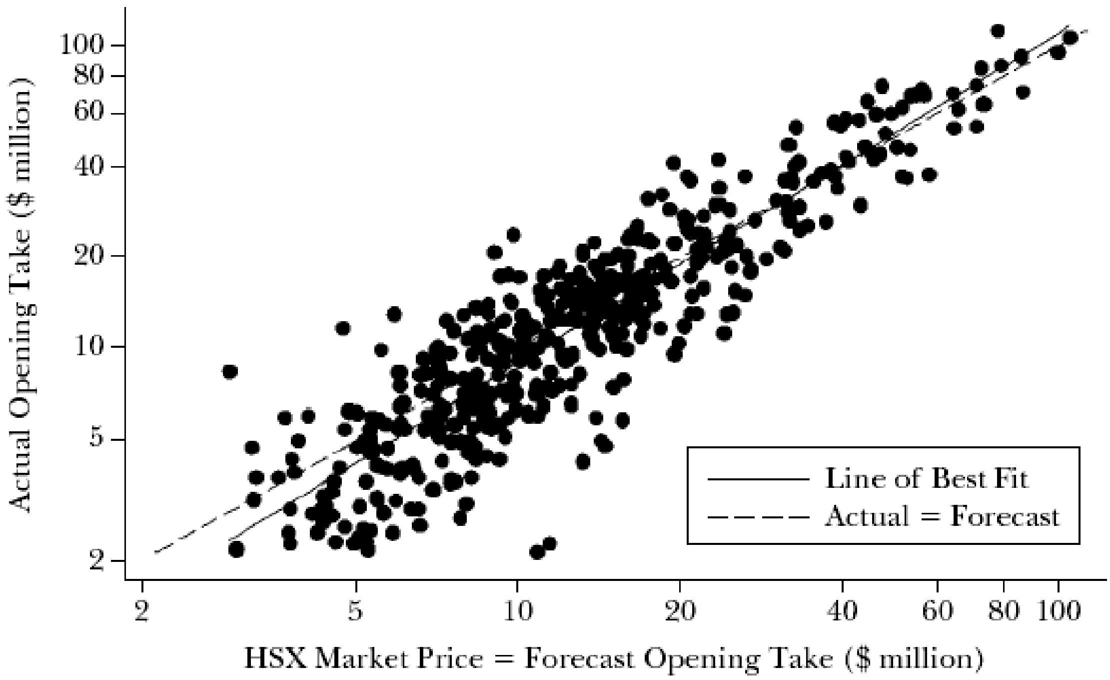

Figure 15: Prediction accuracy of HSX (Wolfers and Zitzewitz, 2004)

Another stream of research proposes the use of prediction markets for long-term forecasting (Hanson, 1992). The Foresight Exchange was already presented as a market for predicting long-term developments in science, technology, and other fields of public interest in Section 2.3.3. Contracts traded in this market range from technical to socio-political issues. Pennock et al. (2001a) show that prices of the Foresight Exchange also correlate well with observed outcome frequencies.

To sum up, prediction markets work well compared to alternative forecasting methods in various fields of application. In addition to these empirical findings on the accuracy of prediction markets in various fields of application, Sunstein (2006) theoretically compared the characteristics of prediction markets to those of the statistical mean of individual judgments and of group judgments generated through deliberation. Sunstein (2006) concludes that prediction markets have substantial advantages to both approaches because deliberation suffers from some serious problems. One of these problems is that group members may not reveal their knowledge due to social pressures. This argumentation is of particular interest since deliberation is widespread and oftentimes considered to be the best way of eliciting information held by groups.

4.1.2. Sports Prediction Markets

Over the past years, prediction markets were also employed in the field of sports forecasting. To test how much extra accuracy can be obtained by using real-money versus play-money prediction markets, Servan-Schreiber et al. (2004) compare the trading prices of the real-money market TradeSports and the play-money market NewsFutures across 208 NFL games. Interestingly, there is no significant difference in the prediction accuracy of play-money versus real-money markets. If two teams are playing against each other, the team with the higher trading price can be considered the favorite. 65.9% of TradeSports' favorite teams won compared to 66.8% of NewsFutures' favorite teams. For both markets, there is a close correspondence between trading prices and the observed frequency of victory in the field (see Figure 16). This shows that the trading prices can be interpreted as probability estimations of the actual outcomes. Rosenbloom and Notz (2006) also find that both markets, TradeSports and NewsFutures, provide accurate probability forecasts.

Moreover, Servan-Schreiber et al. (2004) compare the trading prices of the two prediction markets to the accuracy of predictions from 1,947 individual experts in a popular Internet prediction contest, namely the ProbabilityFootball contest[32]. They find that at the end of the season the markets were ranked 6th (play-money) and 8th (real-money), therewith both falling within the top ten among almost two thousand experts. For comparison, the average expert was ranked 39th and thus still outperforming the

[32] http://ProbabilityFootball.com

majority of individual experts, but not performing as well as the two prediction markets.

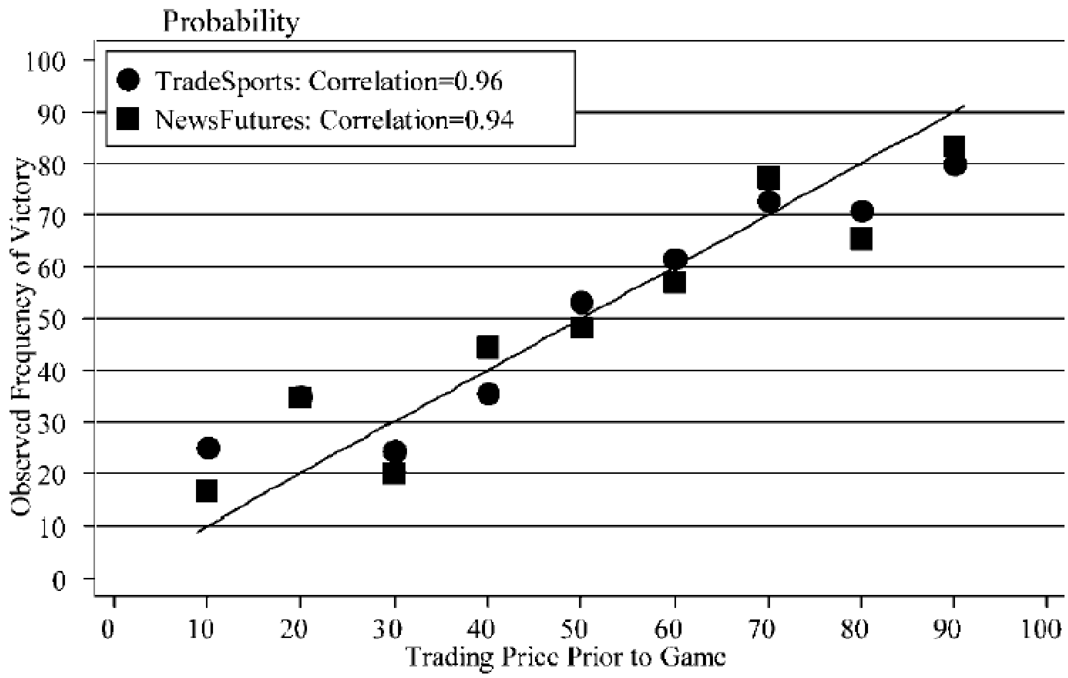

Figure 16: Prediction accuracy for NFL games (Servan-Schreiber et al., 2004)

In a similar study, Chen et al. (2005) find that prediction markets on NFL games offer as accurate predictions as experts' assessments at the same point of time prior to the matches. Furthermore, by analyzing data from 34 soccer prediction markets for the 2002 FIFA World Cup and 18 basketball games from the 2002 USA National Basketball Association (NBA) championship, Debnath et al. (2003) show that on average trading prices approach the actual outcome over time and that information on game events is rather quickly reflected in trading prices.

The first soccer prediction markets to be reported in literature date back to 1994 and 1998 (Schmidt and Werwatz, 2002). In their working paper, Schmidt and Werwatz (2002) analyze a 2000 European Championship market to find out whether prediction markets outperform a random predictor and betting odds across 21 matches. Market-generated probabilities are therefore compared to professional betting odds and a random predictor. The random predictor performed worse than the markets' predictions. Also, relative to the prediction markets, forecasts by expert bookmakers in fixed-odds betting were slightly less accurate.

Overall, there are only few studies on the accuracy of prediction markets in the field of sports forecasting. These studies, though, suggest that markets should also work well in the domain of sports.

4.2. Description of the Data

This section describes the different data sources which are later used to predict the outcome of the 2006 FIFA World Cup and to compare the prediction accuracy of markets to other forecasting methods. The data includes the relevant STOCCER championship and match markets as well as betting odds from two major betting companies, the FIFA world ranking, and a random predictor. The comparison based on this data differs from the study by Schmidt and Werwatz (2002) in several respects.

One of the key features of the soccer prediction markets studied by Schmidt and Werwatz (2002) was the real-money investment which was required: every trader had to deposit a certain amount of money (up to 50€) and thus could suffer losses. As such, these markets were similar to the Iowa Electronic Markets, which have proven to be accurate in the past. In the STOCCER play-money markets, however, traders were not required to make any real-money investments. Traders could therefore neither lose nor win any money by revealing their expectations. Another difference is that the STOCCER prediction markets were more liquid than the markets described by Schmidt and Werwatz (2002). Moreover, in addition to comparing the markets' predictions to betting odds and random predictors as done by Schmidt and Werwatz (2002), the following sections also investigate whether the STOCCER prediction markets outperform forecasts that are based on historic soccer data and to what extent predictions based on different types of contracts diverge.

4.2.1. STOCCER Match Markets

As already described in Section 3.2.1 there were 16 match markets in STOCCER which focused on predicting the outcome of matches in the final rounds. There were three contracts per match. Either one of the two national teams won or there was a draw after the second half. The contract corresponding to the outcome that actually occurred was valued at 10 virtual currency units while the other two contracts became worthless. The matches, the outcome of the matches, and the trading prices of the three possible outcomes are depicted in Table 6.

The trading prices shown in Table 6 are prices of the last trade before kick-off. According to the efficient market hypothesis, these prices incorporate all relevant information available to the traders at this time. For the comparison of forecasting methods in Section 4.3.1, the predicted outcome of a match in case of the match markets is the one with the highest trading price out of the three possible outcomes. In 9 out of the 16 matches, the contract with the highest trading price corresponded to the actual outcome.

Table 6: Trading prices of STOCCER match markets

Match (Team 1 – Team 2)	Last Price Team 1	Draw	Team 2	Result (Team 1 – Team 2)
Germany – Sweden	9.00	0.30	1.60	2-0
Argentina – Mexico	8.28	2.79	1.91	1-1
England – Ecuador	8.75	3.89	2.00	1-0
Portugal – Netherlands	5.40	1.00	4.40	1-0
Italy – Australia	8.90	0.99	1.99	1-0
Switzerland – Ukraine	7.53	1.50	2.40	0-0
Brazil – Ghana	9.50	0.70	0.70	3-0
Spain – France	3.50	1.30	4.99	1-3
Germany – Argentina	6.00	3.75	3.50	1-1
England – Portugal	3.76	2.70	4.05	0-0
Italy – Ukraine	6.70	2.35	1.04	3-0
Brazil – France	6.16	3.22	3.67	0-1
Germany – Italy	5.10	2.28	3.50	0-0
Portugal – France	2.50	3.49	4.92	0-1
Germany – Portugal	5.90	2.50	2.16	3-1
Italy – France	4.50	3.19	3.91	1-1

4.2.2. STOCCER Championship Market

Another set of predictions for all the matches can be derived from the contract prices of the competing teams in the STOCCER championship market, which was also described in more detail in Section 3.2.1. Contracts of all 32 national soccer teams were traded in this market. The matches, the outcome of the matches, and the trading prices of the two teams playing the corresponding match are depicted in Table 7.

Again, the trading prices shown in Table 7 are prices of the last trade before kick-off. These prices should incorporate all relevant information available to the traders at this time.

Table 7: Trading prices of the STOCCER championship market

Match (Team 1 – Team 2)	Last Price Team 1	Last Price Team 2	Result (Team 1 – Team 2)
Germany - Costa Rica	19.99	2.17	4-2
Poland - Ecuador	5.47	2.85	0-2
England - Paraguay	13.48	2.93	1-0
Trinidad & Tobago - Sweden	1.15	7.97	0-0
Argentina - Ivory Coast	16.30	4.30	2-1
Serbia & Montenegro - Netherlands	2.61	11.84	0-1
Mexico - Iran	7.15	2.20	3-1
Angola - Portugal	2.10	7.29	0-1
Australia - Japan	3.26	4.20	3-1
USA - Czech Republic	3.62	8.05	0-3
Italy - Ghana	13.49	1.99	2-0
South Korea - Togo	3.80	1.64	2-1
France - Switzerland	10.31	6.65	0-0
Brazil - Croatia	31.35	4.88	1-0
Spain - Ukraine	8.00	5.19	4-0
Tunisia - Saudi Arabia	3.10	1.43	2-2
Germany - Poland	19.95	2.22	1-0
Ecuador - Costa Rica	5.35	2.00	3-0
England - Trinidad & Tobago	14.20	1.10	2-0
Sweden - Paraguay	6.61	3.51	1-0
Argentina - Serbia & Montenegro	17.05	1.75	6-0
Netherlands - Ivory Coast	11.20	5.20	2-1
Mexico - Angola	7.45	0.65	0-0
Portugal - Iran	7.62	0.31	2-0
Czech Republic - Ghana	12.10	1.25	0-2
Italy - USA	13.40	0.70	1-1
Japan - Croatia	1.40	5.50	0-0
Brazil - Australia	30.94	4.97	2-0
France - South Korea	10.15	4.85	1-1
Togo - Switzerland	0.85	7.45	0-2
Saudi Arabia - Ukraine	0.96	5.18	0-4
Spain - Tunisia	13.75	0.86	3-1
Ecuador - Germany	6.41	20.99	0-3
Costa Rica - Poland	0.04	1.00	1-2
Sweden - England	6.50	13.50	2-2
Paraguay - Trinidad & Tobago	0.03	2.70	2-0
Portugal - Mexico	8.02	5.00	2-1
Iran - Angola	0.06	1.82	1-1
Netherlands - Argentina	11.25	25.10	0-0

Match (Team 1 – Team 2)	Last Price Team 1	Last Price Team 2	Result (Team 1 – Team 2)
Ivory Coast - Serbia & Montenegro	0.06	100.00	3-2
Czech Republic - Italy	7.70	11.20	0-2
Ghana - USA	3.82	2.00	2-1
Japan - Brazil	0.72	29.35	1-4
Croatia - Australia	5.15	4.94	2-2
Saudi Arabia - Spain	0.05	11.55	0-1
Ukraine - Tunisia	6.00	2.30	1-0
Togo - France	0.80	6.50	0-2
Switzerland - South Korea	7.70	4.29	2-0
Germany – Sweden	23.00	5.34	2-0
Argentina – Mexico	28.40	5.04	1-1
England – Ecuador	14.00	5.63	1-0
Portugal – Netherlands	8.37	11.60	1-0
Italy – Australia	18.10	6.20	1-0
Switzerland – Ukraine	13.00	7.18	0-0
Brazil – Ghana	30.20	5.70	3-0
Spain – France	13.95	9.99	1-3
Germany – Argentina	28.45	23.00	1-1
England – Portugal	16.20	16.00	0-0
Italy – Ukraine	19.92	12.85	3-0
Brazil – France	31.01	15.29	0-1
Germany – Italy	41.09	25.65	0-0
Portugal – France	27.00	39.99	0-1
Germany – Portugal	19.79	19.79	3-1
Italy – France	42.00	40.00	1-1

For the analysis in Section 4.3.1, the predicted winner of a match is the team with the higher trading price before kick-off. A draw is predicted whenever the trading prices of two teams are equal. In 38 out of the 64 matches, the team with the higher trading price was the actual winner of the match.

4.2.3. Betting Odds

In fixed-odds betting, one or several professional experts of a betting company set fixed quotes which are usually not adjusted over time (e.g. Forrest et al., 2005). Bettors then accept or reject those bets at some time before the beginning of the respective event. Essentially, in fixed-odds betting information from potentially knowledgeable bettors is not accounted for when determining the odds. Numerous studies have shown that fixed-odds betting markets are efficient (Gandar et al., 1998, Pope and Peel, 1989). For instance, Pope and Peel (1989) develop a linear probability model which incorporates the probabilities of the actual occurrences of the outcomes and the probabilities implicitly quoted by the odd-setters. They then derive several betting strategies and show that no strategy leads to expected positive returns. Nevertheless, some inefficiencies such as the favorite-longshot bias were detected (e.g. Cain et al., 2000, Thaler and Ziemba, 1988). This means that favorites are undervalued and long shots, i.e. outcomes which are very unlikely, are overvalued. For a recent summary of the history of sports wagering see Vlastakis et al. (2006).

In order to avoid losses, betting companies are required to make accurate predictions (Forrest et al., 2005). With large sums of money at stake, the monetary incentive to predict accurately is pronounced and presumably much stronger than in any prediction market since there is no money at stake in play-money markets and usually little money at stake in real-money markets. Forrest et al. (2005, p. 552) emphasize the importance of accurate forecasts for bookmakers in fixed-odds betting markets: “If bets are mispriced, the financial consequences for bookmakers may be serious”. Although a commission fee of 15-25% is usually charged (Woodland and Woodland, 1994) and can palliate possible losses in the short run, under competition, betting companies setting the quotes have a strong incentive to generate accurate quotes. Moreover, one of the bookmakers’ aims is to set the quotes in a way that the bettors’ investments distribute evenly on all three outcomes because the bookmakers do then not take any risk (Schmidt and Werwatz, 2002).

In the following sections, betting odds of two major German sports betting providers, namely ODDSET and wetten.de, are used as a benchmark for the STOCCER prediction markets. ODDSET[33] is Germany's largest betting institution and is run by the state-owned lottery. Wetten.de[34] is a popular sports betting provider that is privately held. Both bookmakers offered fixed quotes which bettors could wager against at the time of the FIFA World Cup 2006. A typical betting screen of wetten.de is depicted in Figure 17.

[33] www.oddset.de
[34] www.wetten.de

Figure 17: Typical screen of a fixed-odd betting site

For each of the 64 World Cup matches, bets could be placed on a win of the first team (1), a draw (0), and a win of the second team (2). All bets are referring to the score after regular playing time. Extra time and penalty shootouts in the final rounds are not considered. Matches that are not decided within regular time are considered a draw. Betting quotes are stated in decimal odds – a bet quoted with 3.5 pays out 3.5 times the wagering amount in case the corresponding event actually occurs. As bookmakers follow a commercial interest and try their best to avoid short-term losses, the odds include a commission fee. This means that wagering the same amount of money on all three possible outcomes would lead to a 15-25% loss. Since soccer is a popular sport in Germany, one can assume that a considerably large amount of money has been betted on outcomes of matches during the FIFA World Cup 2006.

The matches, the outcomes of the matches, and the quotes from wetten.de are depicted in Table 20 (see Appendix A). Respectively, the data from ODDSET is depicted in Table 21 (see Appendix A). For the comparison in Section 4.3.1, the predicted outcome of a match is the one with the lowest quote because according to the quotes this is the most likely outcome. For wetten.de, the outcome with the lowest quote corresponded to the actual outcome of the match in 43 out of the 64 matches. For ODDSET, the actual outcome was predicted for 37 out of the 64 matches.

4.2.4. FIFA Ranking

The FIFA world ranking[35] is a ranking system for men's national soccer teams. The teams of the member nations of the FIFA (Fédération Internationale de Football Association) are ranked according to their match results. The most successful team is ranked highest. In the following, the FIFA world ranking is used as another benchmark since it is based on historic data only. Thus, one can investigate whether the STOCCER prediction markets outperform predictions derived from historic data only and hence do not consider up-to-date information about the current status of the national soccer teams such as players dropping out due to medical reasons or due to disqualification.

The FIFA world ranking from May 2006 which is used as a benchmark in the following takes into account the history of the last eight years before May 2006. The ranking is based on the teams' performance, with more recent and more important matches being weighted more heavily in order to reflect the state of the team. It considers the following factors:

- Outcomes of past matches
- Importance of past matches
- Strength of opponents
- Regional strength
- Results in home and away matches
- Number of goals scored

All international "A" matches are relevant for the calculation of the ranking. For each individual factor, points are assigned which are then aggregated to an index value. In case of most factors complex calculations are used to determine the actual state and strength of the national teams[36].

The matches, the outcomes of the matches, and the ranks of the competing teams in the FIFA world ranking from May 2006 are depicted in Table 22 (see Appendix A). For the analysis in Section 4.3.1, a win is predicted for the team that has the better position

[35] http://www.fifa.com/worldfootball/ranking/

[36] The calculation of the ranking is rather complex. Due to its complexity the calculation procedure was changed in the meanwhile. More information on the calculation of the ranking can be found at http://www.fifa.com/mm/document/fifafacts/rawrank/ip-590_10e_wrpointcalculation_8771.pdf.

in the ranking. This prediction corresponds to the actual outcome for 30 out of the 64 matches.

4.2.5. Random Draws

Forecasts are worthless if they are not better than randomly drawing one of the possible outcomes. Thus, a random predictor is used as another benchmark to evaluate the prediction accuracy of the STOCCER markets. Since one can observe three possible outcomes per match, an uninformed, random guess would correctly predict 33.33% of the matches. Empirical data supports the hypothesis that the three possible outcomes of a match are equally likely to occur (Schmidt and Werwatz, 2002).

4.3. Results

This section compares the prediction accuracy of the STOCCER markets to a random predictor, predictions derived from the FIFA world ranking, and betting odds. First of all, the results of this comparison are given in Section 4.3.1. Since prediction markets should work well if they are efficient, Section 4.3.2 discusses one specific facet of market efficiency. It addresses the question whether pure arbitrage opportunities across contracts existed in the STOCCER championship market. Finally, Section 4.3.3 analyzes whether traders try to exploit the illiquidity of prediction markets by taking on the role of market makers.

4.3.1. Evaluation of the Prediction Accuracy

Predictions based on a random predictor, the FIFA world ranking, and betting odds from two major betting companies are used as benchmarks for the STOCCER prediction markets in order to compare markets to an uninformed guess, to predictions based on historic data only, and to expert predictions by bookmakers. Prediction market prices and thus also the corresponding predictions are, in contrast, driven by the information and the expectations of traders (Spann and Skiera, 2003). Beside historic data, traders also consider current information that is available to them and ongoing developments during the tournament.

In order to compare the prediction accuracy of markets to the other forecasting methods, the hit rate was calculated for each method. The hit rate is the number of correctly predicted matches relative to the total number of predicted matches. How an outcome for a match is predicted in each of the data sets has already been explicated in the last section. Other common evaluation criteria such as the root mean squared error or the mean absolute error for the deviation between the final value of a contract and the last trading price before kick-off cannot be used for comparing the predictions due to the characteristics of the data sets. It is, for instance, impossible to derive probabilities for outcomes of matches from the FIFA world ranking or the trading prices in the championship market. Thus, the hit rate is used as an evaluation criterion which can be employed for all the data sets.

Table 8 compares the hit rate of the different forecasting methods for the whole sample of 64 matches. In case of the STOCCER championship market, a win is predicted for the team with the higher trading price. For the betting odds, the predicted outcome is the one with the lowest quote. The FIFA world ranking predicts a win for the higher-ranked team and in case of the random predictor all three possible outcomes of a match are equally likely to occur.

Table 8: Comparison of prediction accuracy (all matches)

Method	No. Obs.	Hit rate	% improvement[37]	p-value[38]
Championship market	64	59,38%		
Wetten.de odds	64	67,19%	-11,62%	0,203
ODDSET odds	64	57,81%	2,72%	0,799
FIFA world ranking	64	46,88%	26,66%	0,042
Random draw	64	33,33%	78,14%	< 0,001

The comparison of the hit rates of the championship market, the betting odds, the FIFA world ranking, and the random predictor for all 64 matches shows that the championship market indeed yields a higher hit rate than the FIFA world ranking and the random draw model. The difference in the hit rate of the prediction market and these two other forecasting methods is significant in both cases (Pearson's chi-square test, p-value < 0.05)[39]. The predictions can thus be improved when using a prediction

[37] Percentage of improvement of championship market over alternative forecasting method
[38] Chi-square test for difference to hit rate of championship market
[39] For more information on Pearson's chi-square test see e.g. Cowan (1998)

market instead of these two methods. Table 8 shows the percentage of improvement when one replaces the respective alternative method with a prediction market.

With regard to the hit rate, the betting odds from wetten.de and ODDSET perform similarly well as the predictions derived from trading prices before kick-off in the championship market. Wetten.de slightly outperforms the championship market whereas ODDSET performs almost equally well compared to the market. The difference in the hit rate, however, is not significant in both cases. This can be considered as a success for the prediction market because the prediction accuracy obviously is similarly good as in case of betting odds. This is even more astonishing as the market was a play-money market and was also used to predict the course of the entire tournament instead of focusing on the prediction of the outcome of individual matches.

Moreover, the likelihood of draws is systematically underestimated in the championship market. Based on the trading prices in the championship market, a draw would only be predicted if the prices of the competing teams were exactly the same – which is rather unlikely. This also holds for the FIFA world ranking where a draw would only be predicted if two teams were ranked equally.

For this reason, Table 9 compares the prediction accuracy of the various forecasting methods for only those matches out of the total 64 matches which did not end in a draw. In this case, there are only two possible outcomes.

Table 9: Comparison of prediction accuracy (all matches without draws)

Method	No. Obs.	Hit rate	% improvement[40]	p-value[41]
Championship market	47	80,85%		
Wetten.de odds	47	89,36%	-9,52%	0,138
ODDSET odds	47	78,72%	2,71%	0,711
FIFA world ranking	47	63,83%	26,66%	0,003
Random draw	47	50,00%	61,70%	< 0,001

[40] Percentage of improvement of championship market over alternative forecasting method
[41] Chi-square test for difference to hit rate of championship market

The betting odds were adjusted to ignore the probability of a draw by predicting the winner based on which team had the lower odds for it winning the match. However, this does not change the results compared to Table 8. Although again not statistically significant, wetten.de still performs a little better than the championship market while ODDSET is marginally beaten by the market. Also, the championship market still has a much higher hit rate than the FIFA world ranking and the random draw model.

In STOCCER, there were match markets for the 16 matches in the final rounds of the FIFA World Cup 2006. In case of the match markets, the outcome with the highest trading price out of the three possible outcomes is the predicted outcome. Table 10 compares the predictions of these 16 match markets to the predictions of the other forecasting methods.

Table 10: Comparison of prediction accuracy (final rounds)

Method	No. Obs.	Hit rate	% improvement[42]	p-value[43]
Match markets	16	56,25%		
Championship market	16	37,50%	50,00%	0,131
Wetten.de odds	16	43,75%	28,57%	0,313
ODDSET odds	16	43,75%	28,57%	0,313
FIFA ranking	16	25,00%	125,00%	0,012
Random draw	16	33,33%	68,77%	0,044

For the last 16 matches of the tournament, the hit rate of the match markets is significantly higher than the hit rate of the FIFA world ranking and of the random draw model. Interestingly, the hit rate is higher in case of the match markets than it is when predicting a win for the team with the higher trading price in the championship market. One reason for this tendency could again be the fact that the likelihood of draws is underestimated in the championship market. Furthermore, traders in match markets can focus on the outcome of one match at a time instead of trying to predict the course of the entire tournament. In the final rounds, the match markets also seem to outperform the betting odds of wetten.de and ODDSET – although the difference is not statistically

[42] Percentage of improvement of championship market over alternative forecasting method
[43] Chi-square test for difference to hit rate of championship market

significant. Moreover, with only one hit fewer, the prediction accuracy of the championship market is again very close to the prediction accuracy of the betting odds.

Altogether, the STOCCER markets are about as accurate as betting odds and more accurate than the FIFA ranking and a random predictor. At first sight, it is somewhat surprising that the hit rate for the championship market, the betting odds, and the FIFA world ranking is on average lower for the last 16 matches than it is when taking into account all 64 matches. However, this is plausible since it should be easier to predict the outcome of matches at the beginning of the tournament than at the end. At the beginning, there are numerous underdogs and clear favorites whereas towards the end of the tournament the performance of teams will not differ that much. Thus, it is presumably much more demanding to predict the outcome of matches taking place in the last rounds compared to earlier matches.

4.3.2. Arbitrage Opportunities

In case of STOCCER the markets predicted the outcome of the matches quite accurately. Prediction markets should work well if they are efficient. In efficient markets, in turn, one does not expect arbitrage opportunities to be persistent. This section therefore investigates whether pure arbitrage opportunities existed in one specific market, namely the STOCCER championship market. This market was chosen for the following analysis since it was the most liquid market and the only market which was running continuously over a time period of several weeks. Other aspects of market efficiency such as how fast newly arriving information is incorporated into trading prices are not considered here.

In the STOCCER championship market, there are two combinations of trades that can potentially yield arbitrage profits: Firstly, buying all the 32 contracts traded in the market and selling a basic portfolio or, secondly, buying a basic portfolio and selling all the contracts separately in the market. In the first case, one gets paid off on exactly one contract with certainty. If the total of the ask prices on all the contracts is less than 200 currency units at any point in time, an arbitrage opportunity is available. Instead of selling a basic portfolio a trader can also hold the shares until the end. In the second case, the arbitrage opportunity is present if the sum of all the 32 bid prices is more than 200 currency units.

Figure 18 shows the movement of the sum of bid and ask prices in the STOCCER championship market over time. Most of the time the ask prices sum up to more than 200 currency units. Contrariwise, the sum of the bid prices is in the majority of cases lower than 200 currency units. As was already mentioned above, an arbitrage opportunity exists if the sum of bid prices exceeds or the sum of the ask prices falls below 200 currency units. However, extremely small arbitrage opportunities are presumably not of interest for traders because they do not yield any profit worth mentioning in comparison with the effort which is required to trade a portfolio and 32 contracts.

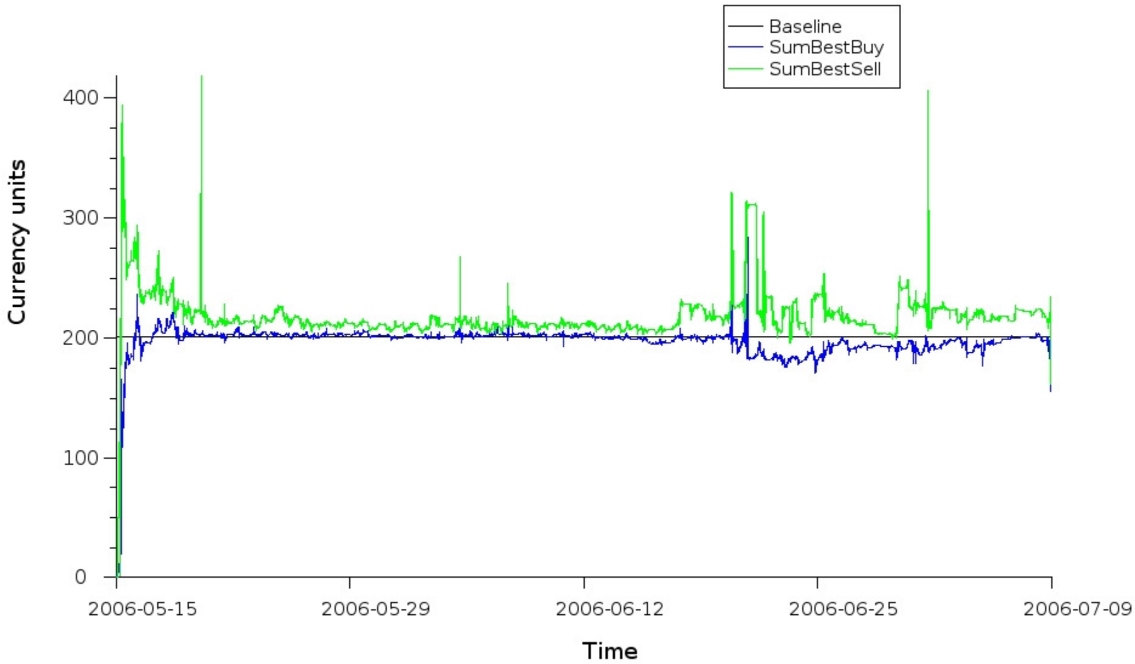

Figure 18: Sum of bid/ask prices in championship market over time

When tolerating arbitrage opportunities of up to one percent of the value of a basic portfolio, i.e. two currency units, there were a total of 229 instances in which an arbitrage opportunity was present between May 15th and July 9th. The arbitrage chances lasted, on average, for about 47 minutes. When tolerating arbitrage opportunities of up to ten percent of the value of a basic portfolio, the number of instances in which an arbitrage opportunity is present declines to seven instances which lasted for 11 minutes on average. Thus, with increasing sums of money at stake the number of arbitrage opportunities declines and substantial arbitrage opportunities are quickly corrected.

Given that trading in this market was relatively thin compared to financial stock markets, it is interesting that the arbitrage opportunities were rather quickly corrected by the traders – provided that a substantial amount of (virtual) money was at stake. All in all, the STOCCER championship market appears to have been efficient in the sense that there were few substantial arbitrage opportunities available by trading basic portfolios or simply holding shares until the outcome was known.

4.3.3. Market-Making Traders

Market liquidity can also become an issue in prediction markets since trading is in many cases relatively thin compared to financial stock markets. If markets are rather illiquid, however, new information is not immediately reflected in trading prices and traders might in consequence lose interest in the markets. One observation worthy of note in case of STOCCER is the emergence of market making traders, i.e. traders who provide liquidity by offering to buy and sell a substantial number of shares of a specific contract at the same time. Market makers add to the liquidity and hope to make profit due to the spread between the buying and selling price.

In the following, the threshold for the number of shares which have to be offered on the buy and sell side at the same time in order to qualify as a market-making trader is 50. Furthermore, taking into account whether the corresponding buy and sell orders were submitted within a given time frame can be seen as an additional constraint. Short time frames imply that traders acted as market makers on purpose. To give an example, it is very unlikely that a trader forgot about a sell order or has completely different information when he submits a buy order for the same contract only a little later.

Table 11 depicts the number of active traders who ever traded a specific contract as well as the number of market-making traders per contract in the STOCCER championship market[44]. On average, there are 622 active traders and 72 market-making traders per contract. The number of market makers decreases if corresponding buy and sell orders have to be submitted within a shorter time frame in order to qualify as a market maker.

[44] This section again relies on data from the championship market since it was the most liquid market and the only market which was running continuously.

Table 11: Number and share of market-making traders per contract

Contract	# Active	#MM	#MM	#MM	#MM (1h)/
Angola	759	67	62	45	5.93%
Argentina	728	82	69	59	8.10%
Australia	665	77	70	54	8.12%
Brazil	765	88	76	56	7.32%
Costa Rica	679	67	55	45	6.63%
Cote d'Ivoire	684	62	54	41	5.99%
Croatia	567	74	64	47	8.29%
Czech Republic	624	70	61	39	6.25%
Ecuador	608	75	65	42	6.91%
England	661	84	69	53	8.02%
France	630	108	97	77	12.22%
Germany	735	102	90	81	11.02%
Ghana	616	81	74	50	8.12%
Iran	628	42	38	25	3.98%
Italy	633	84	72	59	9.32%
Japan	597	58	49	32	5.36%
Korea Republic	547	74	69	47	8.59%
Saudi Arabia	587	55	52	36	6.13%
Mexico	560	72	65	50	8.93%
Netherlands	611	82	71	51	8.35%
Paraguay	570	61	51	36	6.32%
Poland	609	60	51	37	6.08%
Portugal	547	77	64	49	8.96%
Serbia & Montenegro	597	55	49	32	5.36%
Spain	556	82	72	59	10.61%
Sweden	565	77	68	45	7.96%
Switzerland	599	68	54	46	7.68%
Togo	602	49	45	32	5.32%
Trinidad & Tobago	624	67	58	43	6.89%
Tunisia	567	67	57	36	6.35%
Ukraine	571	69	63	54	9.46%
USA	612	69	63	44	7.19%

From now on, this time frame is one hour to be considered a market-making trader. In this case, 7.6 per cent of the active traders are regarded as market makers on average across contracts.

In total, there are 289 different market makers. Some traders are acting as market makers for multiple contracts. Six traders, for instance, qualify as market making traders for more than 25 and up to 31 out of the 32 contracts. Table 12 shows the number of traders who are acting as market makers for multiple contracts. All in all, buying and selling the same contract at the same time seems to be a common trading pattern for some of the traders.

Table 12: Traders acting as market makers for multiple contracts

#Contracts	**1-5**	**6-10**	**11-15**	**16-20**	**21-25**	**> 25**
#MM (1h)	203	42	20	13	5	6

Market-making traders are on at least one side of the trade in 81 per cent of the total contracts traded and account for 85 per cent of the trading volume[45]. The number of trades as well as trading volumes per contract increase with the number of traders who qualify as market makers for a specific contract[46]. Figure 19 shows the correlation between the number of market makers and the number of trades. The correlation coefficient of 0.827 indicates a high correlation between those two numbers[47]. With a correlation coefficient of 0.875, the correlation between the number of market-making traders and trading volumes which is depicted in Figure 20 is similarly high[48].

Hence, both correlation coefficients are high and could reflect the fact that additional market-making traders increase liquidity. However, an alternative explanation could be that the factor which generates trading interest also encourages market makers to trade in the corresponding market.

[45] The market makers' share of trades and trading volume per contract can be found in Table 23 (see Appendix A).

[46] The number of market makers, the number of trades as well as trading volumes per contract can be found in Table 24 (see Appendix A).

[47] Spearman's rank correlation coefficient, p-value < 0.001. For more information on Spearman's rank correlation coefficient see Hotelling and Pabst (1936)

[48] Spearman's rank correlation coefficient, p-value < 0.001

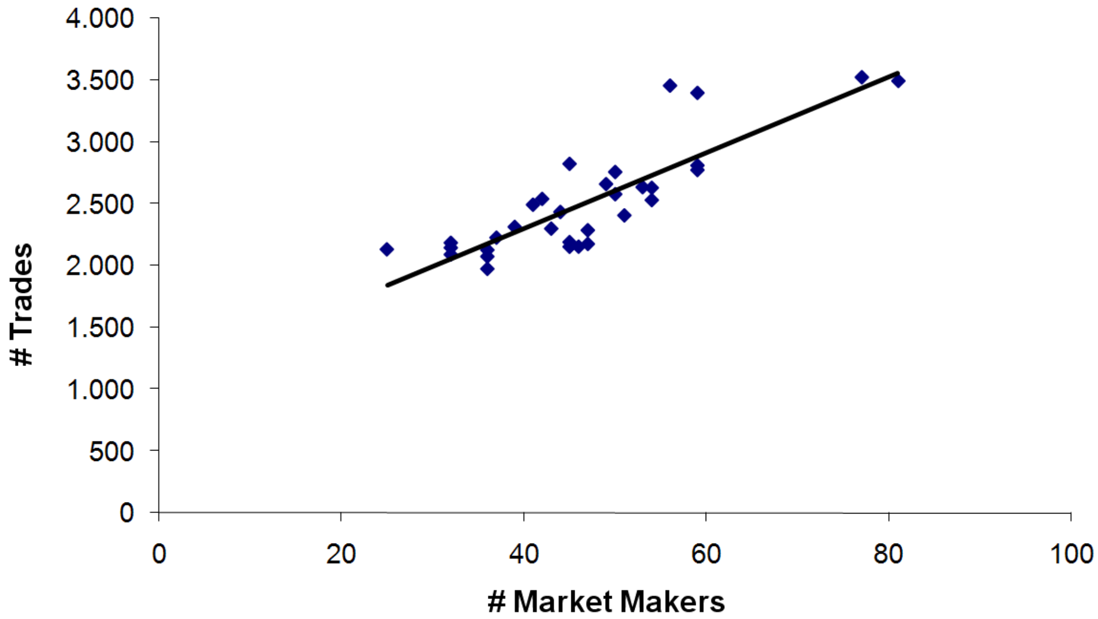

Figure 19: Correlation between number of market makers and number of trades

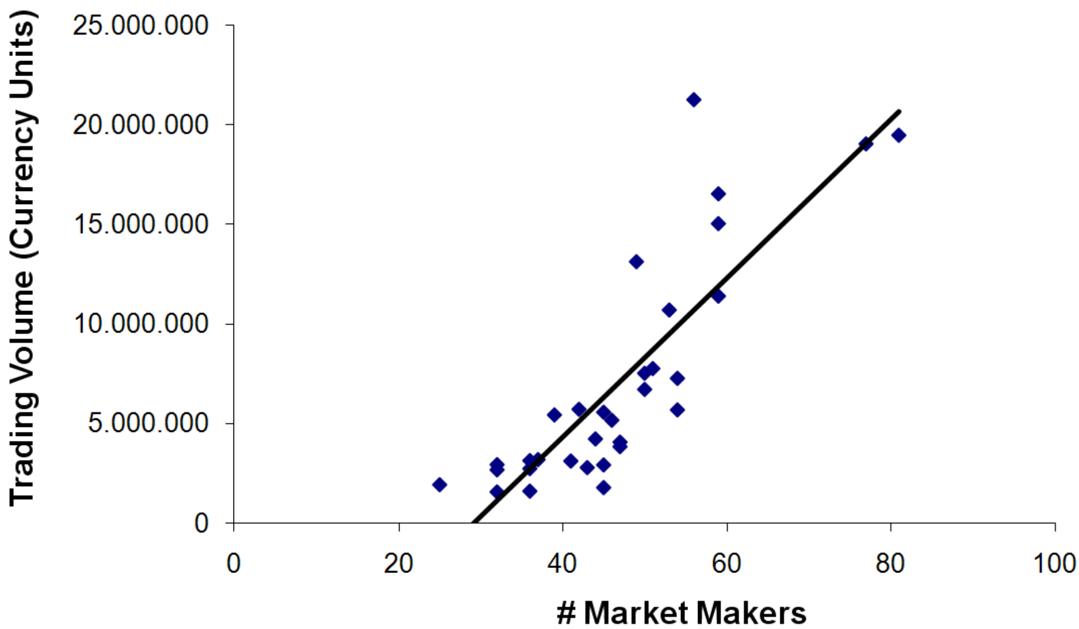

Figure 20: Correlation between number of market makers and trading volume

Without much doubt market makers expect to make profits with their trading strategy of buying and selling specific contracts at the same time. Table 14 shows the market-making as well as the other traders' deposit value, i.e. the sum of the cash and the value of the contracts they hold, at the time when the FIFA World Cup was over and the market had been closed. The average deposit value of market makers is 183,976.52

currency units compared to 135,073.69 currency units for all the remaining traders. The difference between the two groups of traders with regard to the deposit value is significant (Mann-Whitney U test, p-value = 0.003)[49]. Market makers thus are more successful than the remaining traders with respect to their deposit value.

Table 13: Trading activity and trading success of market makers[50]

	MM	Non-MM	p-value[51]
Mean number of trades	413.62 (719.56)	43.21 (61.55)	< 0.001
Mean deposit value	183,976.52 (165,738.49)	135,073.69 (62,443.13)	0.003

As shown in Table 13, market-making traders are also trading a lot more than other traders. On average, market makers trade about 414 times whereas other traders only make about 43 trades. Again, the difference in the number of trades is significant. Market makers obviously try to profit from illiquidity. Thus, they play an important role in prediction markets by providing liquidity and consequently allowing for continuous trading.

4.4. Discussion of Results

The results reported in Section 4.3.1 provide evidence that the STOCCER prediction markets in fact outperformed predictions derived from the FIFA world ranking, i.e. historic data, as well as a random predictor in terms of prediction accuracy. What is more, the differences between prediction markets and betting odds were not statistically significant with respect to hit rates. Overall, quotes from wetten.de tend to be a little more accurate than quotes from ODDSET. Predictions based on the championship market were about as accurate as the betting odds although the probability of draws is underestimated by this market and the focus was not on predicting the outcome of individual matches. As a consequence, the match markets should perform a little better. Differences in the hit rate between the match markets and the betting odds are, however, also not statistically significant.

[49] For more information on the Mann-Whitney U test see Mann and Whitney (1947)
[50] The numbers in parentheses are standard deviations.
[51] The p-values are obtained from a Mann-Whitney U test.

Yet, when keeping in mind that betting odds have shown to be extremely good predictors and that similar forecasting methods are mostly non-existent in other fields of application beyond sports forecasting, the results confirm that prediction markets indeed are a very promising forecasting method.

In play-money prediction markets traders cannot suffer any losses whereas professional bookmakers depend on accurate predictions due to the high monetary investments that are at stake. In the end, betting companies with inaccurate quotes would not survive. By demonstrating the competitiveness of the STOCCER prediction markets compared to sports betting odds, the results align with those attained by Schmidt and Werwatz (2002) where markets even slightly outperformed betting odds. However, their markets used real money as an incentive and traders hence were punished financially in case of a poor performance. Play-money prediction markets as the ones analyzed here, however, are much easier to set up and to operate than real-money prediction markets due to legal and technical reasons. It is thus crucial to find out whether prediction accuracy decreases when using play money instead of real money.

In addition to the comparison of prediction markets and alternative forecasting methods, the analysis of arbitrage opportunities in Section 4.3.2 also gives evidence that the markets are performing well. The championship market appears to be efficient in the sense that there are few substantial pure arbitrage opportunities available. Besides, Section 4.3.3 demonstrates the importance of market-making traders who provide liquidity and consequently allowing for continuous trading in otherwise rather illiquid markets. They seem to play a central role in prediction markets since they are on the buy or sell side in a large proportion of all trades.

4.5. Summary

This chapter provided evidence of markets' prediction accuracy. Earlier empirical research was used to demonstrate the predictive power of markets relative to traditional forecasting methods such as expert opinions or polls in general as well as for sports forecasting in particular. Data collected from the play-money prediction market STOCCER was then employed to compare the prediction accuracy of sports prediction markets to alternative forecasting methods. After a short description of the data sets, the prediction accuracy of the STOCCER markets was analyzed by comparing the predictions to a random predictor, predictions derived from the FIFA world ranking, and betting odds from two German betting companies.

The results showed that the play-money prediction market STOCCER for the FIFA World Cup 2006 was about as accurate as betting odds. Betting odds, in turn, are known to be very accurate predictors. Moreover, the markets clearly outperformed predictions based on the FIFA world ranking as well as the random predictions. The chapter contributes to the literature by providing the first empirical comparison of play-money prediction markets and predictions based on historic data or betting odds in the field of sports forecasting.

An analysis of the championship market documented that prediction markets also appear to be efficient in the sense that there are few substantial arbitrage opportunities available. Furthermore, it was shown that market markers play an important role in prediction markets. They serve as liquidity providers and allow for continuous trading.

The chapter concluded with a discussion of the results regarding the prediction accuracy of the STOCCER markets. In consideration of these results and related empirical research, one can ascertain that prediction markets in many cases perform at least as well as alternative forecasting methods. Quite often they even outperform well-established methods such as polls in the field of political elections.

5. Incentive Schemes for Play-Money Prediction Markets

In comparison to most alternative forecasting methods, prediction markets provide a rather sophisticated incentive scheme as traders can be rewarded based on their performance which is a good indicator for the quality of their contributions. So far, market operators have employed various kinds of incentive schemes in order to motivate people to participate in such markets and to reveal their expectations. Typical incentive schemes include prizes for the top performers of a market, lotteries among all traders, rankings published on the World Wide Web or even real-money exchanges where traders have to invest some of their own money. In real-money prediction markets you "put your money where your mouth is" (Hanson, 1990a). This is advantageous with respect to opinion weighting since traders will put more money on predictions they are more confident about. Furthermore, it is reasonable to expect that the prospect of gaining or losing money will motivate traders to seek accurate information. However, real-money prediction markets are illegal or at least highly regulated in most countries. Even in countries that offer betting licenses to prediction market operators, setting up a real-money market incurs huge technical and regulatory costs. System failures or attempts to defraud can become business-critical. Real-money markets thus are hard to realize. Moreover, potential traders might be unwilling to invest their own money in prediction markets.

Previous research in the field of prediction markets as well as the results described in the preceding chapter have shown that play-money as well as real-money markets can predict future events at a remarkable degree of accuracy. In play-money prediction markets, though, well-designed incentive schemes are oftentimes needed to encourage participation and information revelation. The popularity of several play-money markets demonstrates that such incentives can motivate intense trading (Robinson, 2001). Presumably, the embodiment of the incentive scheme has a huge impact on market efficiency and the accuracy of predictions.

In this chapter, three different monetary incentive schemes for play-money prediction markets are compared with regard to their impact on the accuracy of predictions. In order to do so, predictions from three groups of traders corresponding to three treatments with different incentive schemes are studied in a field experiment. Subjects of the first group received a fixed amount of money, subjects of the second group were paid according to their ordinal rank, and the third group the subjects' payment depended linearly on their deposit value in the prediction market. Studying these incentive schemes is of special interest when traders need to be paid for taking part in a prediction market, e.g. in the case of an internal market for company-specific predictions. In such a market it is improbable that employees risk some of their own money in order to generate better company forecasts. Based on the results of the field experiment, advice on engineering incentive schemes for prediction markets is given.

The remainder of the chapter is structured as follows. Section 5.1 discusses prior studies on the comparison of real-money and play-money prediction markets as well as related work on incentives schemes in the field of experimental economics. Section 5.2 then describes the field experiment that was conducted during the FIFA World Cup 2006 in addition to the public STOCCER markets described in Chapter 3. Subsequently, Section 5.3 presents the results of the field experiment concerning the impact of the three above-mentioned incentive schemes on the accuracy of predictions. Section 5.4 discusses these results and gives an outlook on possible implications for designing incentive schemes of public and intra-enterprise prediction markets. Finally, Section 5.5 briefly summarizes the main findings of the chapter.

5.1. Related Work

Despite the importance of designing suitable incentive schemes, there are so far only two research papers studying different incentive schemes for prediction markets and these compare real-money and play-money markets. The results of both studies are described in Section 5.1.1. Related experimental studies on financial incentive schemes beyond the field of prediction markets are presented in Section 5.1.2.

5.1.1. Real-money vs. Play-money Prediction Markets

Servan-Schreiber et al. (2004) compare predictions of the real-money market TradeSports to those of the play-money market NewsFutures regarding American Football outcomes of 208 matches during the 2003-2004 NFL season in order to examine how much extra prediction accuracy can be obtained by using real money instead of play money. Both markets among others trade contracts which are valued at 100 if a team wins and 0 if it loses. The real-money market TradeSports requires an initial investment when joining the market and charges a transaction fee on each trade. In contrast, registration in NewFutures is free and each trader initially receives a fixed amount of play money. Very successful traders can use their play money to bid on real prizes. The number of traders per contract under study was known for NewsFutures but not for TradeSports. Servan-Schreiber et al. (2004) assumed that this number "is of the same order of magnitude" in both markets.

In their study, Servan-Schreiber et al. (2004) find that there is no statistically significant difference between the real-money market TradeSports and the play-money market NewsFutures. "Both types of markets exhibited significant predictive powers" and had almost the same prediction accuracy (Servan-Schreiber et al., 2004). Their trading prices correlated well with actual outcome frequencies. This result raises the question how the draw between play-money and real-money prediction markets can be explained. First of all, there is supposedly intrinsic interest in NFL football and no reason not to trade truthfully except team biases. Additionally, the weights given to the traders' opinions reflect the amount of money they are willing to put on their predictions in real-money markets. This is most likely affected by their wealth levels rather than by their trading success in prediction markets. In play-money markets, in contrast, the traders' wealth depends on a history of accurate predictions and opinion weights should thus be more efficient.

In a second study on the comparison of play-money and real-money prediction markets, Rosenbloom and Notz (2006) analyzed sports events such as baseball games, basketball games, hockey games, tennis matches, or golf tournaments and also predictions on the direction of financial markets and political events such as whether John Edwards will be chosen as the vice presidential candidate in the 2004 United States presidential elections. Overall, the correlation between TradeSports probabilities and NewsFutures probabilities was 0.955. In case of team sports such as NFL games, Rosenbloom and Notz (2006) produced conclusions consistent with those from Servan-Schreiber et al. (2004), i.e. they did not find any statistically significant difference between real-money and play-money markets for NFL games in particular and team sports in general. In spite of this, Rosenbloom and Notz (2006) found TradeSports to be slightly but significantly more accurate than NewsFutures for all the other events. They conclude dthat predicting the outcome of matches in case of team sports might be different from other events. Rosenbloom and Notz (2006) speculated that traders are getting cues from other sources such as betting odds in case of team sports. This might influence their assessment of the probabilities.

In consideration of both studies, the impact of real money vs. play money on the accuracy of predictions is not completely understood and clarified. Moreover, there exists far more than one design alternative only for play-money markets – and also for real-money markets. The strength of both studies is the large data set from real-world online experiments that both studies rely on. However, both studies do not consider any other differences apart from the use of real money or play money in their comparison of the two markets. Although the markets they compare are quite similar, they are far from identical. Without doubt, a key difference between the two markets is that one uses real money while the other does not. But how did other aspects influence prediction accuracy?

It remains an open question how, for example, the number of traders and their trading activity influence the market and thus also the accuracy of predictions. This seems to be an interesting question since the number of traders per contract was not available to Servan-Schreiber et al. (2004) in case of TradeSports. Rosenbloom and Notz (2006) report that in their case there were far more traders in NewsFutures than in TradeSports. The number of contracts traded in NewsFutures ranged from 95 to 157,891 contracts with a mean of 7,600 contracts. In TradeSports it ranged from 1 contract to 21,771 contracts with a mean of 201 contracts. Thus, the real-money market TradeSports performed better despite the lower volume.

Furthermore, TradeSports also levies a small fee on each transaction. How does this impact trading behavior and thus also trading prices? The two markets – TradeSports and NewsFutures – were not identical in all respects and it therefore remains an open question how other factors have influenced the results described by Servan-Schreiber et al. (2004) and by Rosenbloom and Notz (2006).

5.1.2. Experimental Studies on Monetary Incentives

Beyond prediction markets, financial incentives are also employed to improve the performance of employees in companies or subjects in economic experiments. Monetary incentives are supposed to motivate people to exert additional effort and should thus improve task performance. Many experimental economists would in all probability insist that monetary risk is required to obtain valid conclusions about economic behavior (Servan-Schreiber et al., 2004). In a survey of 31 experimental studies, Smith and Walker (1993) found that increased monetary incentives indeed bring the behavior closer to the predictions of economic theory with rational agents. Monetary incentives also reduce the variance of the data around the outcome predicted by rational models. This can presumably be explained by the fact that subjects in an experiment balance monetary incentives against decision costs and thus rather deviate from rational predictions in case of lower incentives. Regarding monetary incentives in companies, recent evidence suggests that employees perform better when using performance-based compensation schemes for which the payment is closely related to performance (Prendergast, 1999). Although incentive plans are sometimes criticized because they do not alter the employees' attitudes and might make them less enthusiastic about their work (Kohn, 1993).

On the other hand, there is empirical evidence indicating that monetary and performance-related incentives do not necessarily increase performance. In several experiments Gneezy and Rustichini (2000), for example, find that offering money does not always improve performance. For small monetary incentives they observed a decrease in performance compared to treatments with zero compensation. However, if money was offered, a larger amount of money resulted in a higher performance. Possible explanations could be that subjects follow social norms or are intrinsically motivated independent of any monetary incentive. The level of intrinsic motivation most likely depends on the task at hand and on individual interests. One of the experiments conducted by Gneezy and Rustichini (2000) was a donation experiment where subjects collected donations from the public. If subjects have to conduct such a useful task, one would expect a rather high level of intrinsic motivation. Gneezy and Rustichini (2000) speculate that the introduction of a monetary incentive displaces intrinsic motivation.

It is indeed a frequently discussed issue in psychology and education that monetary incentives decrease intrinsic motivation. So far, a consensus has not been reached in experimental psychology. However, a recent meta-analysis of incentives and intrinsic motivation suggests that, in general, monetary incentives are not harmful to subjects' motivation to perform a task (Cameron et al., 2001). Cameron et al. (2001) state that intrinsic motivation increases or does not differ from a control group without any monetary incentive if monetary incentives are performance-based.

Overall, meta-studies such as the one by Camerer and Hogarth (1999) show that the presence and amount of monetary incentives seem to affect average performance in many tasks. Even in cases where incentives do not change average behavior substantively, the variance of responses often decreases. However, the effects of monetary incentives on performance are mixed and complicated. On the one hand, they depend on the tasks to be performed (Camerer and Hogarth, 1999). Incentives affect the performance of subjects in particular if increased effort improves performance. On the other hand, the effects of monetary incentives depend on the type of incentive schemes which is employed. In a review of 131 laboratory experiments, Bonner et al. (2000) studied the relation between the type of incentive scheme and subjects' task performance. On the whole, their review reveals that monetary incentives improve performance in about half of the experiments. Furthermore, Bonner et al. (2000) find that not all incentive schemes elicit the same level of effort. Piece-rate schemes, for instance, have a higher likelihood of positive incentive effects than tournament schemes which are in turn followed by fixed-rate schemes.

Earlier research in the field of labor contracting also confirms that employees perform better in case of performance-based compensation schemes compared to fixed rates (Prendergast, 1999). Concerning the comparison of tournaments and piece rates, Lazear and Rosen (1981) theoretically demonstrate that rank-order tournaments are often efficient and yield an allocation of resources identical to that generated by piece rates. Under some circumstances, risk-averse employees should actually prefer to be paid on the basis of rank, i.e. according to their ordinal rank in the organization rather than their output level. Which compensation scheme is preferred by the employees depends on the utility function and the amount of luck involved (Lazear and Rosen, 1981). In an experimental study on the comparison of tournament and piece rates, Bull et al. (1987) find that the mean effort levels chosen by the subjects converged to their theoretical equilibrium levels for both compensation schemes. However, there was a large variance of behavior across identical tournaments in case of the rank-order tournament whereas the variance was quite small in case of the piece-rate scheme. Bull et al. (1987) attribute this difference to the fact that rank-order tournaments require strategic behavior whereas piece rates simply induce maximizing behavior.

Beyond these thoughts on how to design suitable incentive schemes, Bewley (1995) notes that managers in real-life know from experience that one should not rely on monetary incentives alone to motivate employees. Employees – especially those who have contact with the public – should for instance be happy, and happiness cannot be achieved with monetary incentives alone.

5.2. A Field Experiment on Monetary Incentives in Prediction Markets

Studying the impact of different monetary incentive schemes on the prediction accuracy of markets is an open and interesting object of investigation. A field experiment was conducted to analyze alternative monetary incentive schemes in a prediction market. These incentive schemes could for instance be used in internal markets for company-specific predictions. Traders can then be rewarded for joining a market and contributing to accurate forecasts.

This section describes the setup of a field experiment which was conducted in parallel to the public STOCCER markets during the FIFA World Cup 2006. Firstly, the basic setup of the field experiment is presented. Secondly, the three monetary incentive schemes that were compared with regard to their prediction accuracy are described and it is explained why exactly those three incentive schemes were chosen for the study. Thirdly, based on previous research the expected results of the experiment are discussed.

5.2.1. Basic Setup

The underlying events used for the field experiment were the outcomes of soccer matches. Similar to the match markets in the public STOCCER markets (see Section

3.2.1) there were 20 markets for the last 20 matches of the FIFA World Cup 2006. Contracts traded in the markets were the possible outcomes of all the matches. There were three possible outcomes for every match – either one of the two national soccer teams won or there was a draw after the second half, i.e. at the end of the regular playing time. The third contract "draw" was traded although there were no draws possible in 16 out of the 20 matches. The reason was that the outcome of overtimes and penalty shootouts was considered to be more or less unpredictable. The contract corresponding to the event that actually occurred during the World Cup was valued at 100 currency units after the match; the other two assets were worthless.

In total, 60 undergraduate students from the Universität Karlsruhe (TH), Germany, were taking part in the field experiment in June and July 2006. The operational principle of prediction markets was briefly explained in a lecture and students could then volunteer for the field experiment. After registering for the experiment they received subsequent instructions via e-mail[52]. Moreover, the students were asked to complete a short pre-experiment questionnaire in order to collect demographic data and information about the students' risk attitude. All the markets opened two days before the corresponding match and closed at the end of the match[53]. Traders were able to buy and sell basic portfolios comprising the three contracts traded in a market at 100 currency units at any time. This way, contracts were placed into circulation. The trading mechanism was a standard continuous double auction (CDA) with an open order book and limit orders. Short selling was not permitted. The trading software used for the experiment was the same one as in case of the public STOCCER markets (cp. Section 3.3).

5.2.2. Incentive Schemes

The 60 students were randomly assigned to three groups of 20 students each. At the end of the FIFA World Cup 2006 the traders were paid in real money according to their group's incentive scheme. This allows for studying the impact of three different

[52] See Instruction 1 in Appendix B (in German) for the instructions which were sent to the subjects via e-mail when trading started for the first time. Depending on the incentive scheme the text of this e-mail slightly varied.

[53] See Instruction 2 in Appendix B (in German) for the instructions which were sent to the subjects via e-mail when the first match markets were launched. Subsequently, the subjects were informed via e-mail whenever new match markets were launched.

monetary incentive schemes by comparing the prediction accuracy of the three groups of traders, corresponding to three treatments with different incentive schemes. The subjects of the first group were paid a fixed amount of 50 Euro irrespective of how successful they traded in the markets (from now on referred to as fixed payment, FP). In the second group, individuals were paid according to their ordinal rank (rank-order tournament, RO). The trader ranked first within the group was paid 500 Euro, the second 300 Euro, and the third 200 Euro. All the other traders in this group did not receive any payment at all. Although the average payment is also 50 Euro per person, in this case, few traders win big prizes. Subjects in the third group were promised what was called a performance-compatible payment, also with an average amount of 50 Euro (deposit value, DV). Performance-compatible means that the payment linearly depended on the traders' success, i.e. the deposit value in the prediction market (deposit value divided by 10.000), and was therefore directly influenced by every transaction a trader carried out.

These three incentive schemes were chosen for the field experiment because they are closely related – although they admittedly are not exactly the same – to incentives that can nowadays typically be observed in public as well as corporate prediction markets. In case of public markets, there are usually markets without any payment or prizes to win, markets with rank-order tournaments, and real-money markets. Similarly, comparing the three monetary incentive schemes is also of interest for operators of internal markets for company-specific predictions. Companies are oftentimes willing to reward their employees' effort and so far used various incentives such as rankings demonstrating the expertise of successful traders, rank-order tournaments with big winners, and real-money markets where the employees' investments are subsidized by the company. These incentive schemes are again similar to the ones investigated in this field experiment and consequently the question arises which incentive scheme is the most suitable.

For every group, the 20 markets on 20 soccer matches of the FIFA World Cup were run separately, i.e. the same market existed three times. Aside from the difference in the incentive schemes, the market environment was identical across groups. This facilitates a more reliable test of the effect of incentives in prediction markets than has been reported in any of the related literature. Since subjects who did not trade at all should also not receive any payment, a relatively small minimum trading volume was imposed on all traders. The minimum weekly trading volume corresponded to 5 Euro in real money, i.e. 10 per cent of the initial deposit value. The weekly trading volume was displayed in the trading screen and subjects consequently always knew how much they had to trade in order to reach the minimum trading volume. Especially in the case of the fixed payment group subjects might otherwise have considered not trading at all or simply could have forgotten to participate in the online experiment.

5.2.3. Expected Results

Based on earlier research on monetary incentives it is to be expected that the performance-compatible group performs best and the fixed payment group performs worst in terms of prediction accuracy. In the following the intuition behind these expectations is explained.

On the one hand, no extrinsic motivation is given to subjects of the fixed payment group to reveal their expectations or to be among the top performers of the group. Basically, there is no incentive for them to trade more than the minimum required trading volume per week. One would thus expect a rather low trading activity compared to the other incentive schemes. On the other hand, one should not forget about the traders' intrinsic motivation and also their interest in soccer. Traders receiving a fixed payment independent of their performance should not display a reduction in intrinsic motivation compared to unrewarded groups of traders (Gneezy and Rustichini, 2000, Cameron et al., 2001). They may also consider it a duty to perform well in exchange for receiving the payment of 50 Euro. Furthermore, traders do not risk any money and risk aversion thus does not come into play. Nevertheless, it can be suspected that the fixed payment scheme performs worse than the other two incentive schemes since it is known that the presence of performance-based monetary incentives does enhance average performance in many tasks (Camerer and Hogarth, 1999, Prendergast, 1999).

Members of the third group receive a performance-compatible payment, meaning that every transaction directly influences the payment they receive. Traders are paid according to their individual output level and should consequently be motivated to try their best. Due to their risk aversion, traders probably try to avoid losing money and consider very carefully what and how to trade. One can for this reason expect that traders from this group trade less and at slightly lower prices. Their increased effort, however, should improve their performance (Camerer and Hogarth, 1999) and therefore also prediction accuracy. In short, traders with the incentive scheme DV have to “put their money where their mouth is” (Hanson, 1999) and consequently predictions are expected to be rather accurate. In contrast to e.g. rewarding corporate executives, where it is difficult to observe an individual’s output, this is straightforward in case of traders’ performance and hence not a downside of the performance-compatible payment.

For the rank-order tournament one can expect a result somewhere in between the other two groups. On the one hand, traders have a strong incentive to be among top three traders of their group because they will not receive any payment otherwise. This should lead to a rather high trading activity. Moreover, rank-order tournaments have also been considered as a promising payment scheme in other contexts such as labor contracting (Bull et al., 1987, Lazear and Rosen, 1981). On the other hand, the rank-order tournament provides an incentive to take higher risks compared to traders receiving the performance-compatible payment. Strategic behavior comes into play because the margin of winning does not affect payments (Lazear and Rosen, 1981). Also, traders might start betting on unlikely events because they consider this the best or maybe even only way to outperform their competitors. At least, traders falling behind are likely to take risky strategies to catch up with competing traders (Prendergast, 1999). Also, traders could stop trading as soon as they assume that they do not have any chance of becoming one of the top three traders of their group.

All in all, one could expect that the performance-compatible incentive scheme outperforms the rank-order tournament which in turn does better than the fixed payment. This would also be consistent with findings from laboratory studies which examined the effect of different incentive schemes on performance. In a review on comparable incentive schemes, Bonner et al. (2000) conclude that piece-rate schemes have a higher likelihood of positive incentive effects than tournament schemes which are in turn followed by fixed-rate schemes.

5.3. Experimental Results

This section describes the results from the field experiment on monetary incentives. Firstly, the trading activity of the three treatments is compared in order to find out how the different incentive schemes affect trading activity. Secondly, the distribution of trading prices in the three treatments is analyzed. Finally, the impact of the three different incentive schemes on the accuracy of predictions which are derived from trading prices is examined.

5.3.1. Trading Activity

As was already described in Section 5.2.3, one could probably influence the level of trading in a prediction market by choosing a certain incentive scheme. In case of a fixed payment there is no monetary incentive to trade more than the minimum trading volume whereas a competitive incentive scheme such as the rank-order tournament should stimulate trading. Table 14 shows the total and mean number of trades as well as the standard deviation in the three treatments of the field experiment.

Table 14: Trading activity in the three treatments

Treatment	# trades (total)	# trades (mean)	# trades (std dev)
FP (fixed payment)	1520	76	69.08
RO (rank-order tournament)	962	48.1	42.58
DV (deposit value)	1319	65.95	47.74

Perhaps somewhat surprisingly, with a total of 1,520 the number of trades is highest in case of the treatment with the fixed payment and lowest in case of the rank-order tournament with a total of 962 trades. In the third treatment in which payments are linearly based on the traders' success, the number of trades lies between the other two treatments. Relative to the treatments with performance-based incentive schemes (RO and DV) the trading activity is higher than expected in the group with a fixed payment. The differences in trading activity between the three groups, however, are not statistically significant (Kruskal-Wallis test, p-value = 0.355)[54, 55]. Despite the relatively high trading activity in case of the FP treatment, there was not a single trade in four markets. In the RO treatment, there were still two markets with no trading activity. This is of course undesirable because it is then impossible to derive any predictions from trading prices. The only treatment with trading activity in all markets was the DV treatment.

5.3.2. Trading Prices

In total, every group traded 60 contracts in 20 different markets. Figure 21 illustrates how many contracts were traded within certain price ranges in each of the three treatments. The prices under examination here are the last trading prices before the corresponding match started. Contracts are grouped into five price ranges and, for each treatment, the share of contracts with trading prices in each of the price ranges is depicted. The very first column, for example, shows that before the match started 32% of the contracts were traded at prices between 0 and 20 virtual currency units in the first treatment with a fixed payment. Accordingly, in the RO treatment 19% of the contracts were traded within this price range.

When comparing the three treatments one can see that a relatively high number of contracts were traded at prices between 60 and 100 currency units in the rank-order tournament treatment. Moreover, a relatively small number of contracts were traded at prices between 0 and 20 currency units in this treatment. Subjects are obviously willing to take some risk in treatment with the rank-order tournament and buy contracts even at rather high prices. In case the trading prices are good predictors the likelihood of the underlying events should be similarly high as the prices.

Subjects in the performance-compatible payment group, in contrast, do not trade any contract at a price between 80 and 100 currency units and almost no contract in the price range from 60 to 80. Obviously, traders with the payment scheme DV are

---

[54] The null hypothesis of the Kruskal-Wallis test states that there is no difference between the mean trading activities of the groups. The null hypothesis cannot be rejected here. For more information on the Kruskal-Wallis test see Kruskal and Wallis (1952).

[55] Although the Kolmogorov-Smirnov test shows that distributions in each of the groups are normal, an analysis of variance cannot be used in this case because the variance of the data in the groups is not the same. The Bartlett's test was used to test for equal variances. For more information on the Bartlett's test see Bartlett (1937).

unwilling to take the risk of buying contracts at such a high price although there is no reason why their expectations about specific outcomes of the matches should differ from the traders' expectations in the other two treatments. At the other extreme, 52% of the contracts are traded for less than 20 currency units in the DV treatment.

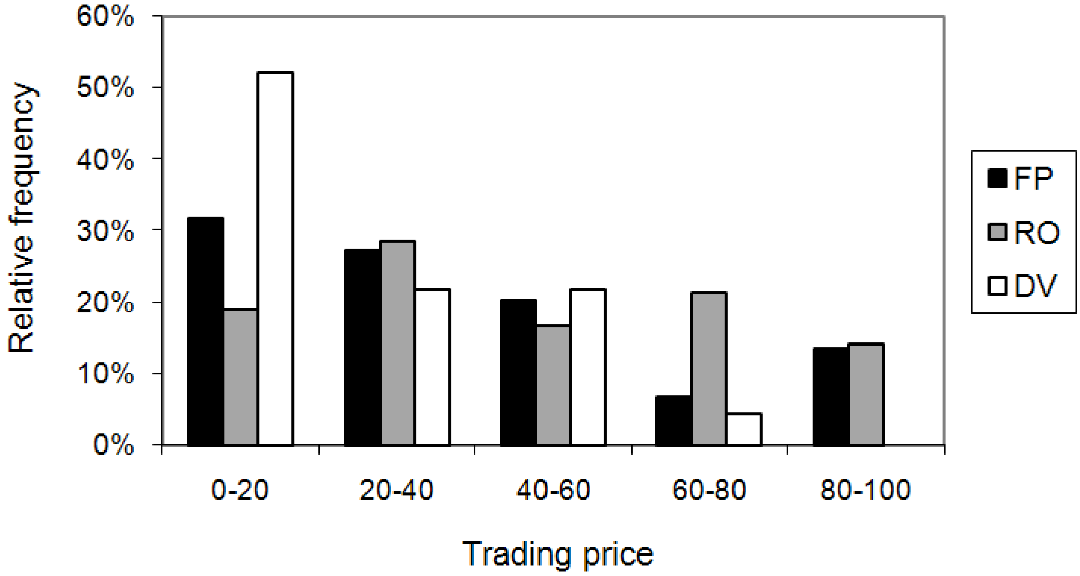

Figure 21: Distribution of trading prices in the three treatments[56]

On average, trading prices for the same matches are lowest in the DV treatment and highest in the RO treatment. One possible explanation for the cautious behavior of traders in the third treatment could be their risk aversion. Due to their risk aversion, traders seem to trade contracts at lower prices compared to the other two treatments. Obviously, they are unwilling to buy contracts at prices similar to the ones in the other treatments and at the same time are willing to sell contracts at rather low prices. Traders in the RO treatment, however, are willing to take some risk in order to outperform the competing subjects of their group. The FP treatment does not impose any monetary risk at all and risk aversion thus should not matter. The following section discusses how this trading behavior impacts the prediction accuracy of the three treatments.

[56] The exact shares of contracts traded within the five price ranges in each of the treatments are given in Table 25 of Appendix B.

5.3.3. Prediction Accuracy

Overall, 35% of the contracts with the highest trading price out of the three contracts per match actually corresponded to the observed outcome in case of the fixed payment. This can also be referred to as hit rate of the markets as it was defined earlier in this work for the match markets of STOCCER which were open to the public (see Section 4.3.1). The average pre-game trading price of the contract corresponding to the outcome was 40.83 virtual currency units. In the rank-order tournament, the most likely outcome according to the trading prices actually occurred in 45% of the cases and the average pre-game trading price of the contract corresponding to the outcome was 51.65 currency units. Finally, in case of the performance-compatible payment, the most likely outcome according to the trading prices occurred in merely 20% of the cases and the average pre-game trading price of the contract corresponding to the outcome was 26.64 currency units. When interpreting the trading prices as probabilities the third group predicted the outcome of a match even worse than the treatment with a fixed payment. The rank-order tournament, in contrast, seems to work quite well with regard to the hit rate and average pre-game trading price. This is rather surprising and contrary to the expected results discussed in Section 5.2.3. However, the differences between the average pre-game trading prices of the three treatments are not statistically significant (Kruskal-Wallis test, p-value = 0.156)[57]. Concerning the hit rate, there can only be found a statistically significant difference between the RO and the DV treatment (Pearson's chi-square test, p-value = 0.024)[58, 59].

Section 5.3.2 already described that trading prices seemed to be rather low in case of the performance-compatible payment compared to the other treatments. This can also be seen when calculating the sum of the three contract prices corresponding to the three possible outcomes of a match. These prices should sum up to about 100 virtual currency units since the probability that one of the three events occurs is 100%. In case of the performance-related incentive scheme the average price of such a so called basic

[57] The null hypothesis of the Kruskal-Wallis test cannot be rejected here and differences between the trading prices are thus not statistically significant.

[58] For more information on Pearson's chi-square test see e.g. Cowan (1998).

[59] Although there was no trading activity in 4 markets in case of the FP treatment and in 2 markets in case of the RO treatment, the hit rate was calculated as the number of correctly predicted matches relative to the total number of matches. The hit rate of those two treatments would otherwise be a little higher. Nevertheless, this is not desirable since markets with no trades at all are also not useful for making predictions about the outcome of matches.

portfolio is only 53.30 virtual currency units while it is indeed very close to 100 in the other two treatments (97.72 in the FP treatment and 102.83 in the RO treatment). This is surprising because there is an arbitrage opportunity in case of such a deviation of the sum of the three contract prices from 100. Traders should buy all three contracts in a market and hold them or sell a basic portfolio since they get paid off on exactly one contract with certainty. But a thorough analysis of incoming and executed orders shows that it was impossible to buy all three contracts in a market at the same time for a sum of prices below 100 currency units. Traders as a consequence could not make use of arbitrage opportunities because the markets were not liquid enough. This also explains why the average pre-game trading price is extremely low in case of the DV group.

To analyze the correspondence between trading prices and outcome frequencies in more detail, the data was sorted into buckets by assigning all of the contracts to one of five price ranges according to their pre-game trading price. Figure 22 plots the relative frequency of outcome against the trading prices observed before the corresponding match started.

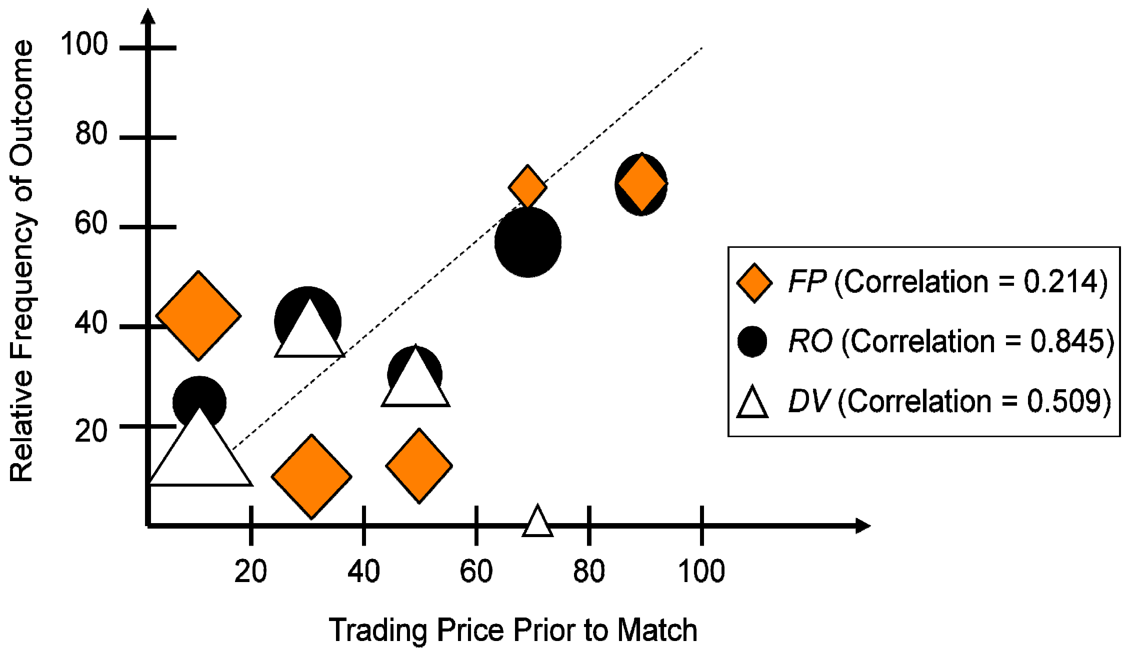

Figure 22: Market forecast probability and actual probability[60]

[60] The relative frequencies of outcome of contracts traded within the five price ranges in each of the treatments are given in Table 26 of Appendix B.

If the markets are efficient, a plot of trading prices vs. observed outcome frequencies should approximate the 45-degree line which represents perfect accuracy. One should thus observe that contracts traded, for example, at a price of 30 currency units correspond to the actual outcome with a probability of 30% on average. The size of the circles, diamonds, and triangles indicates how many trading prices fall into the corresponding price range in case of the different incentive schemes. The larger a circle, diamonds, or triangle is, the more contracts were assigned to this price range.

A first glance at Figure 22 already shows that the trading prices and outcome frequencies seem to correspond rather well in case of the rank-order tournament. The correlation between the relative frequency of outcome and the trading prices serves as an indicator for the accuracy of predictions[61]. For the rank-order tournament, the correlation coefficient is 0.845 which indicates a high correlation between outcome frequencies and trading prices. While there still is a medium correlation of 0.509 in case of the DV group, the correlation is not statistically significant for the predictions from the FP group[62]. Thus, trading prices from the RO group reach the highest correlation with outcome frequencies compared to the other two incentive schemes. Once again, the rank-order tournament seems to outperform the other incentive schemes. In contrast to the expected results discussed in Section 5.2.3, the prediction accuracy here is found to be better in case of the rank-order tournament than in case of the payment based linearly on the trading success in the DV treatment. The FP incentive scheme performs very poor as the correlation between trading prices and outcome frequency did not reach significance.

As was already discussed earlier, on average the sum of the three trading prices corresponding to the three possible outcomes of a match was only 53.30 virtual currency units in case of performance-compatible incentive scheme. Due to the low trading prices in the DV treatment there is no triangle in the price range between 80 and 100 currency units of Figure 22. This might also explain why the prediction accuracy of the treatment with the rank-order tournament is higher. When dividing all the trading prices by the average price of a basic portfolio, in the DV treatment, the correlation

[61] Spearman's rank correlation coefficient is employed to measure the correlation. For more information on this correlation coefficient see Hotelling and Pabst (1936).
[62] p-value < 0.001 for RO and DV; p-value $= 0.082$ in case of FP

coefficient between the relative frequency of outcome and the trading prices after all increases to 0.653[63]. Still, the correlation coefficient is higher in the RO treatment without any need for normalization. This also makes the interpretation of trading prices as probabilities much easier in the RO treatment.

5.4. Discussion of Results

One can only speculate about possible reasons for this result, i.e. in particular the good performance of the rank-order tournament. Traders are obviously not only driven by monetary incentives since they do not stop trading as soon as they reach the minimum weekly trading volume in the FP treatment. Also, in case of the rank-order tournament, traders continue to trade even if winning becomes extremely unlikely for them. This explains why even the markets of the FP group work to some extent. Nevertheless, there was no trading activity for four matches and also no significant correlation between trading prices and outcome frequencies in case of the FP treatment. A fixed payment consequently does not seem to be a well-suited incentive scheme to remunerate traders in a play-money prediction market.

Still, intrinsic motivation does not explain the higher prediction accuracy of the RO treatment compared to the DV treatment since there is no obvious reason why intrinsic motivation should be different in these treatments. Both incentive schemes are performance-based but differ with respect to the accuracy of predictions. The traders' risk aversion could be one reason for the good performance of the rank-order tournament relative to the payment which depends linearly on the traders' success.

Before the field experiment on monetary incentives started, a lottery choice experiment as known from Holt and Laury (2002) was conducted in order to measure the traders' degree of risk aversion. Similar to Figure 23, subjects were presented a menu of choices which permits measurement of the degree of risk aversion[64]. The probabilities were explained in terms of throws of a ten-sided dice. The amounts of money, however, were fifty times the ones shown in Figure 23. The choices thus involved large cash prizes that were paid to the subjects. The payoffs for Option A are less variable than the payoffs of the risky Option B. When the probability of the high-payoff outcome

[63] Spearman's rank correlation coefficient, p-value < 0.001
[64] The screen which was presented to subjects can be found in Figure 30 of Appendix B (in German).

increases enough subjects should cross over from Option A to Option B. A risk-neutral subject would choose Option A four times before switching to Option B.

Option A	Option B	Expected payoff difference
1/10 of $2.00, 9/10 of $1.60	1/10 of $3.85, 9/10 of $0.10	$1.17
2/10 of $2.00, 8/10 of $1.60	2/10 of $3.85, 8/10 of $0.10	$0.83
3/10 of $2.00, 7/10 of $1.60	3/10 of $3.85, 7/10 of $0.10	$0.50
4/10 of $2.00, 6/10 of $1.60	4/10 of $3.85, 6/10 of $0.10	$0.16
5/10 of $2.00, 5/10 of $1.60	5/10 of $3.85, 5/10 of $0.10	−$0.18
6/10 of $2.00, 4/10 of $1.60	6/10 of $3.85, 4/10 of $0.10	−$0.51
7/10 of $2.00, 3/10 of $1.60	7/10 of $3.85, 3/10 of $0.10	−$0.85
8/10 of $2.00, 2/10 of $1.60	8/10 of $3.85, 2/10 of $0.10	−$1.18
9/10 of $2.00, 1/10 of $1.60	9/10 of $3.85, 1/10 of $0.10	−$1.52
10/10 of $2.00, 0/10 of $1.60	10/10 of $3.85, 0/10 of $0.10	−$1.85

Figure 23: Ten paired lottery-choice decisions (Holt and Laury, 2002)

50 out of the 60 subjects from the field experiment also participated in the lottery choice experiment. Only 7 subjects ever switched back from B to A. Figure 24 depicts the average proportion of safe choices in the experiment as well as the risk neutral prediction for each of the ten decisions. One can see that the series of choice frequencies lies to the right of the risk neutral prediction. Across the three groups, nearly 75% of the subjects chose more than four safe choices and thus exhibited risk aversion. These results are in line with those reported in the literature (Holt and Laury, 2002, Harrison et al., 2007, Holt and Laury, 2005).

In case of the fixed payment, traders can neither win nor lose money, so they just play for fun and their risk aversion should not matter. Moreover, traders will take quite a lot of risk in the rank-order tournament because they have to be among the top performers within their group to receive the relatively large cash prize. Thus, the incentives override risk aversion. Only in case of the performance-compatible incentive scheme, traders receive an endowment of 50 Euro and could potentially lose money with every trade they make. As a result, buyers are obviously extremely cautious and not willing to spend too much money on any contract. But why are sellers willing to give up contracts at prices below their average worth? Subjects had to trade in order to reach the minimum transaction volume. Once sellers have started to partially sell their basic portfolios they are probably willing to sell at rather low prices to avoid the risk of holding contracts of an event that does in the end not occur. Average trading prices are thus much lower than in case of the DV treatment than in the two other treatments. Evidently, the performance-compatible payment scheme is less suitable to reveal the traders' expectations about the likelihood of future events than the rank-order tournament.

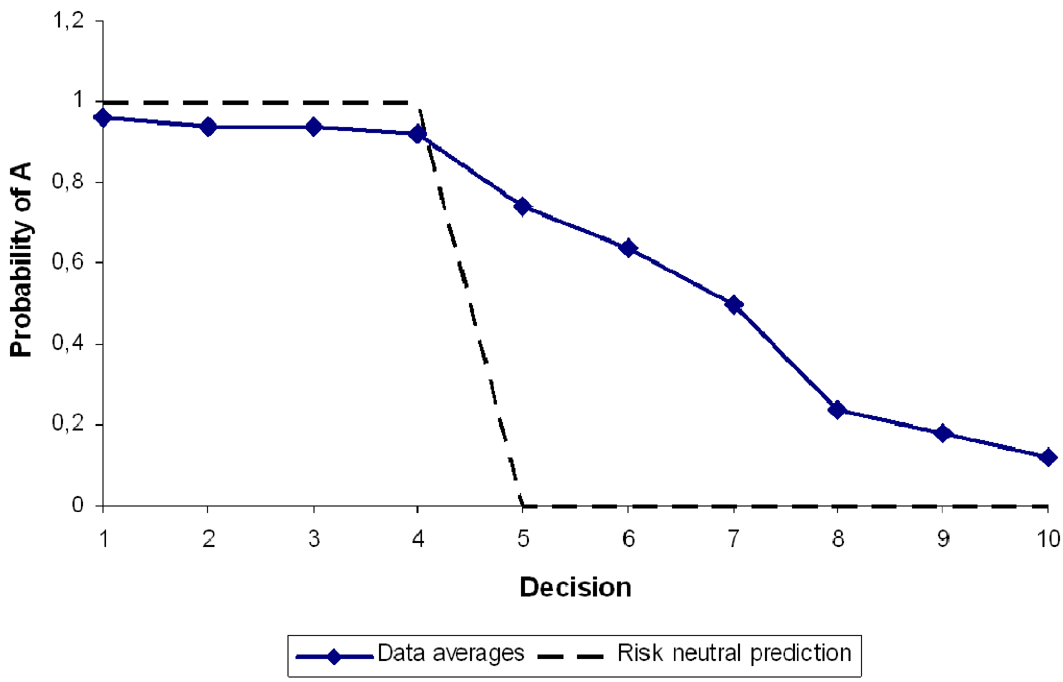

Figure 24: Proportion of safe choices in each decision

But what are the implications for designing incentive schemes of future prediction markets? Out of the three incentive schemes under examination in the field experiment prediction market operators should choose the rank-order tournament when, for example, setting up an internal market for company-specific predictions in which employees are to be rewarded for trading. Besides, performance-compatible payment schemes are somewhat similar to real-money markets. But is it now possible to draw the conclusion that play-money markets e.g. with prizes for the top performers will outperform real-money markets although the latter raise numerous legal and technical difficulties? One should rather be careful when answering this question based on the results of the field experiment because the situation might be somewhat different in prediction markets that are open to the public. In this case, there is a self-selection of traders and it is thus reasonable to expect that many traders in a public real-money market are risk-seeking. In such a situation a performance-compatible payment scheme might potentially produce much better predictions than in the case of the field experiment which is discussed here.

5.5. Summary

This chapter analyzed the impact of different incentive schemes on the accuracy of prediction markets. The only two articles dealing with incentives in the field of prediction markets compare play-money to real-money markets while in the field experiment different incentive schemes of play-money markets were compared. Subsequent to a brief discussion of previous experimental studies on monetary incentives beyond prediction markets, the field experiment studying three commonly-used incentive schemes of play-money prediction markets was presented.

The results show that rank-order tournaments are a suitable incentive scheme in case of risk-averse traders. A lottery choice experiment was utilized to measure the traders' degree of risk aversion. The competitive environment in the corresponding treatment overrides risk aversion and in doing so leads to the best results in terms of prediction accuracy. Due to the risk aversion average trading prices were by far too low in case of the performance-compatible incentive schemes. The fixed payment scheme was also found to be ill-suited since there was no trading activity for several events and also no significant correlation between trading prices and outcome frequencies.

The chapter concluded with a discussion of the results. Those are highly relevant for the question of how to remunerate traders in internal corporate prediction markets but cannot be directly transferred to public real-money prediction markets.

6. Traders' Biases in Prediction Markets

According to the early work by Hayek (1945) and Fama's (1970b) efficient market hypothesis, market prices aggregate and reveal the information traders have. Research in behavioral economics and behavioral finance, however, provides evidence of anomalies in individual behavior. It has been demonstrated that individuals exhibit substantial information processing or judgment biases. Markets which require probabilistic calculations and forecasts of future outcomes are particularly challenging with regard to the traders' information processing capabilities. Traders' biases may thus also affect their trading behavior in prediction markets and in doing so influence predictions based on trading prices. In the first place the question arises whether such biases can actually be observed in prediction markets. If so, it is interesting to study whether markets still work well even if traders do not behave as economic theory assumes.

In the field of political stock markets, Forsythe et al. (1992) for the first time demonstrated that traders are buying and selling contracts of US presidential candidates in a manner which is correlated with their preferences, i.e. supporters of a candidate buy more contracts of this candidate than they sell. This is contradictory to the assumption that rational traders should not trade according to their individual preferences but according to the expected election outcome. Their preferences, however, seem to affect their expectations and traders might unconsciously support their preferred candidate or party. Forsythe et al. (1992) attribute the observed biases to failures in the traders' information-processing capabilities. However, manipulation should be considered as an alternative explanation for the traders' behavior in political stock markets.

So far, empirical results on price manipulation in prediction markets are mixed. In an experimental study, Hanson et al. (2006) find that manipulators are unable to distort prediction accuracy. In contrast to standard prediction markets, subjects in this experiment knew that there were manipulators in the market trying to bull the market. It therefore remains an open question whether the experimental results also hold for prediction markets in general. Hanson (2006) even argues that manipulative trading should usually improve prediction accuracy. Earlier research on manipulation which has been conducted in the field of political stock markets, on the other hand, provides evidence of manipulation. Hansen et al. (2004) discuss the effect of manipulation under the preconditions of indecisive voters and mass media coverage. Under such circumstances, traders might try to influence voters via the predictions which are published in the media. Indeed, Hansen et al. (2004) conclude that political stock markets are vulnerable to manipulation and that trading prices can be manipulated effectively. In a market on the Swedish referendum about joining the European Union, Bohm and Sonnegard (1999) also find that it is possible for a group of traders to distort trading prices at least for a certain period of time.

Overall, the impact of manipulation on the performance of prediction markets is without doubt an open question (Wolfers and Zitzewitz, 2005). Attempts of manipulation could also explain why traders in political stock markets are buying and selling contracts in a manner which is correlated with their preferences. In political elections, traders might try to influence other voters and hence also the outcome of the election. As a consequence, it appears reasonable to study the impact of traders' biases on their trading behavior in a field of application where traders cannot influence the outcome of the corresponding event.

Sports tournaments are such a domain. Traders in STOCCER, for example, in all likelihood cannot influence the outcome of soccer matches or the performance of their national soccer team. This chapter examines whether traders in the sports prediction market STOCCER exhibit any systematic biases resulting from their nationality. It studies the impact of the traders' nationality on their holdings and their trading behavior. If trading is correlated with preferences, traders should buy more and sell fewer contracts of their national team than other traders.

The remainder of this chapter is structured as follows. Section 6.1 outlines related work on biases in financial and prediction markets. Subsequently, Section 6.2 describes the STOCCER data which is used to study the correlation between the traders' nationality and their trading behavior. The results of the study are presented in Section 6.3. Firstly, the correlation between nationality and shareholdings is examined. Secondly, the impact of the traders' nationality on their trading behavior is analyzed. Thirdly, differences in trading prices of prediction markets which focus on specific target groups are studied. Section 6.4 discusses the results and their implications for the selection of traders in prediction markets. Finally, Section 6.5 briefly summarizes the main findings of this chapter.

6.1. Related Work

Biases such as the home bias which describes the tendency of investors to allocate a large fraction of their portfolio to domestic assets are well-known phenomena in financial markets. Section 6.1.1 addresses earlier research on the home bias since it is closely related to the bias studied in this chapter. The home bias in financial markets deals with the impact of investors' nationality on asset holdings whereas the study discussed in this chapter examines how the traders' nationality influences shareholdings and trading behavior in prediction markets. Earlier research on biases in prediction markets is presented in Section 6.1.2. So far, prediction market researchers focus on the favorite-longshot bias and partisanship in political stock markets.

6.1.1. Home Bias in Financial Markets

Investing abroad may improve an investor's risk-return portfolio profile because foreign assets do not always move together with domestic assets. Usually, there is a relatively high degree of correlation within an economy (Levy and Sarnat, 1970). Depending on the correlation of returns across different countries, investors in consequence may benefit from an internationally diversified portfolio. Although this has been known for decades, there is strong evidence that investors allocate only a small fraction of their portfolio to foreign investments (e.g. French and Poterba, 1991, Cooper and Kaplanis, 1994, Tesar and Werner, 1995, Kang and Stulz, 1997). Investors in the United States, for example, allocate only about eight percent of their holdings to foreign assets although the optimal weight is about 40 percent (Pástor, 2000)[65]. Actually, "investors should put much more of their wealth into foreign assets" (Glassmann and Riddick, 2001, p. 35).

Nevertheless, underweighting of foreign assets may be due to rational reasons. Several possible explanations for this underdiversification, which is often referred to as home

[65] The sample period ended in 1996. A more recent publication states that investors in the United States in the meanwhile allocate about 12 percent of their holdings to foreign assets. See(Ahearne et al., 2004)

bias, have been discussed in the academic literature. First, institutional factors such as taxes or greater transaction costs for foreign assets may reduce returns from investing abroad and thus make home assets more attractive (Glassmann and Riddick, 2001, French and Poterba, 1991). However, obstacles to international portfolio investment have decreased dramatically over the years because of, for instance, international tax accords (Kang and Stulz, 1997). At the same time, the home bias has also decreased substantially. Nevertheless, it still remains high (Ahearne et al., 2004).

Institutional factors are of second-order importance in the meanwhile. Thus, other explanations for the home bias have been put forth. The imperfect diversification could as well be caused by information asymmetries and differences in the investors' expectations. French and Poterba (1991) argue that investors expect returns in their domestic asset market to be higher than returns in other markets. Shiller et al. (1991) report survey data which indicates that domestic investors are more optimistic about domestic market returns than foreign investors are. Moreover, investors may attribute a higher risk to foreign investments because they know more about the domestic market and try to bypass political risks which are associated with foreign investments. If the perceived risk of foreign assets increases they become less desirable, thus generating a home bias (Glassmann and Riddick, 2001).

Information asymmetries could also explain the findings of Kang and Stulz (1997) who studied foreign equity ownership in Japan. They observed that foreign investors in Japan did not hold the market portfolio. In fact, investors preferred to hold shares of large and well-known manufacturing companies as well as companies with good accounting performance. Information asymmetries are most likely smaller for these companies compared to small and unknown companies. Large companies, for instance, are more likely to sell their products abroad and are thus known by consumers and potential shareholders in foreign countries. Merton (1987) argues that investors hold shares they know about. This may explain why investors did not hold the Japanese market portfolio. Due to this bias, investors' return in Japan was more volatile compared to holding the market portfolio (Kang and Stulz, 1997).

Across countries, information asymmetries which result from low credibility and poor financial information in some countries can be overcome, for instance, by cross-listing the shares of a company in countries with reputable accounting standards and regulatory environments (Ahearne et al., 2004). Even within countries, then again, investors prefer to invest in local companies (Coval and Moskowitz, 1999). Asymmetric information between local and non-local investors could also explain this preference for geographically close-by investments. After all, local investors may obtain important information about a company from employees, suppliers, and the local media. This can be considered an information advantage compared to other investors.

Overall, most of the explanations of the home bias are based on compelling intuitions but it is an empirical question to determine which explain observed behavior. None of them individually has succeeded in resolving the home bias puzzle (Kang and Stulz, 1997). The home bias can most likely only be explained by a combination of the above-mentioned factors.

While institutional factors are irrelevant when studying traders' behavior in a specific prediction market, information asymmetries as well as differences in traders' expectations might indeed play a role in prediction markets. Within the scope of the study discussed in this chapter, the traders' nationality may influence their standard of knowledge about national soccer teams and their expectations.

6.1.2. Biases in Prediction Markets

Research on biases in the field of prediction markets focuses on the favorite-longshot bias and partisanship in political stock markets. The favorite-longshot bias is well-known from racetrack betting data (e.g. Thaler and Ziemba, 1988, Ali, 1977, Hausch et al., 1981). According to this bias, betting odds provide biased estimates of probabilities. Bettors tend to overvalue longshots and undervalue favorites in betting markets. For this reason, betting on favorites yields higher returns than betting on longshots. As shown by betting data from horse races between 1992 and 2001 in Figure 25, the loss of betting on odds of 100/1 is about 61 percent while betting the favorites yields losses of only around 5.5 percent (Wolfers and Zitzewitz, 2006b).

Although betting markets are generally efficient, bettors' misperceptions of probability drive the favorite-longshot bias (Snowberg and Wolfers, 2007). The favorite-longshot bias may thus be considered the result of market inefficiency (Woodland and

Woodland, 1994). A similar mispricing has also been found in prediction markets. Wolfers and Zitzewitz (2004), for instance, study financial variables which are traded on Tradesports prediction markets. They find that extremely unlikely outcomes are relatively overpriced on TradeSports. Moreover, Leigh et al. (2007) provide evidence of a modest favorite-longshot bias in political stock markets. Overall, the favorite-longshot bias is well-documented in betting as well as prediction markets. Since this chapter examines whether traders in the STOCCER prediction market exhibit any systematic biases resulting from their nationality, the favorite-longshot bias is not studied in more detail in the following.

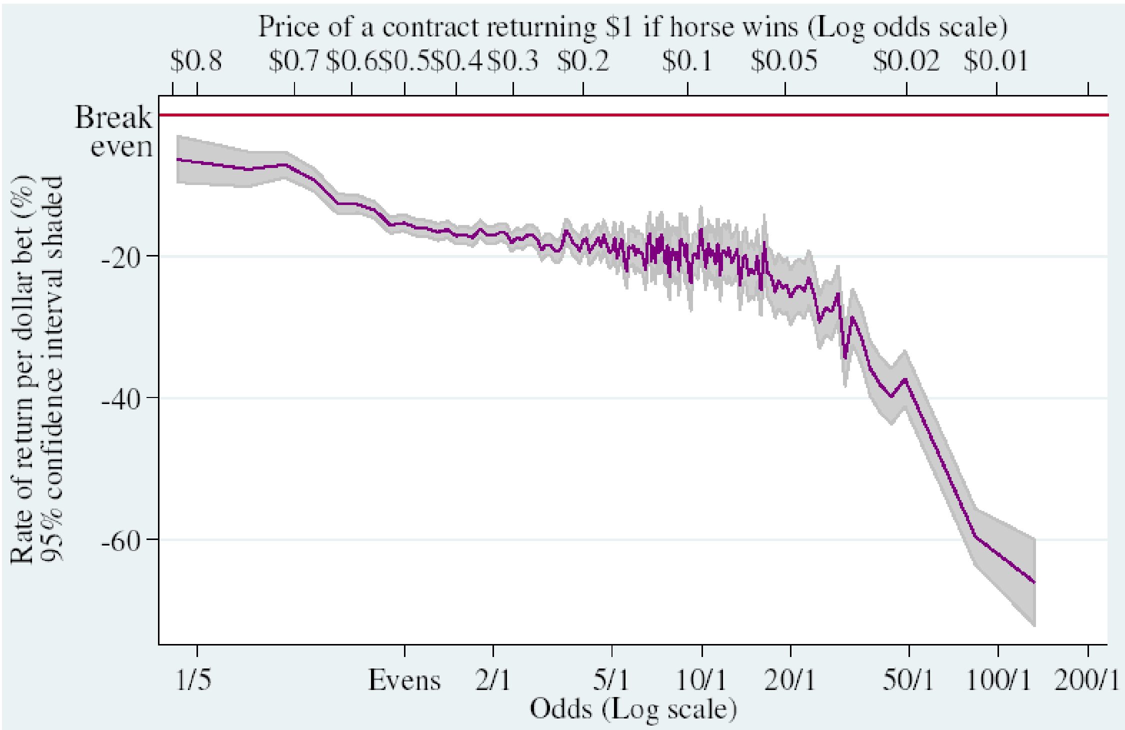

Figure 25: Rate of return at different odds (Wolfers and Zitzewitz, 2006b)

Another behavioral bias observed in prediction markets reflects the tendency of traders to trade according to their desires. Forsythe et al. (1992) demonstrated that traders are buying and selling contracts in a political stock market in a manner which is correlated with their party identification. This is also a common phenomenon beyond prediction markets. Political scientists have observed a tendency to overestimate the preferred candidate's or party's chances of victory (e.g. Bartels, 1987).

Most likely traders' preferences over parties and candidates influence their perception of reality as well as the interpretation of news and information about the likelihood of the outcome occurring. Furthermore, Table 15 demonstrates that individuals overestimate the extent to which their views are representative for all voters. Respondents were asked for which candidate they intend to vote and which candidate they expect to win. The supporters of a candidate are most of the time quite convinced that their candidate is going to win the election. A typical example was the US presidential election in 1980. While 87 percent of the Democrats expected a Democrat to win, merely 19.6 percent of the Republicans expected a Democrat to win.

Table 15: Preferences and expectation in elections (Forsythe et al., 1992)

Year	Democrat/Republican	Percentage of respondents intending to vote Democratic who expect Democrat to win	Percentage of respondents intending to vote Republican who expect Democrat to win
1988	Dukakis/Bush	51.7	5.8
1984	Mondale/Reagan	28.8	1.0
1980	Carter/Reagan	87.0	19.6
1976	Carter/Ford	84.2	19.4
1972	McGovern/Nixon	24.7	0.4
1968	Humphrey/Nixon	62.5	4.6
1964	Johnson/Goldwater	98.6	69.5
1960	Kennedy/Nixon	78.4	15.8
1956	Stevenson/Eisenhower	54.6	2.4
1952	Stevenson/Eisenhower	81.4	14.1

Forsythe et al. (1992) provide evidence that biases affected trading behavior on average by matching individual trading data to political preferences. Supporters of a candidate buy more contracts of the candidate than they sell. This is also referred to as the wishful thinking effect (Forsythe et al., 1999). Despite these judgment biases, the market worked extremely well. Forsythe et al. (1992) explain the success of the market with the so called marginal trader hypothesis. According to this hypothesis, average traders are biased but prices are determined by marginal traders. Marginal traders here are defined as traders who submit limit orders at prices close to trading prices. These traders invested more and traded more actively. Indeed, Forsythe et al. (1992) find that marginal traders did not suffer from judgment bias in their trades. They are presumably motivated by profits rather than partisanship. Nevertheless, there is not direct test for this assumption.

Forsythe et al. (1999) reproduce the wishful thinking effect which was found in political stock markets in laboratory markets. Most probably, there is a tendency to overestimate the probability of desirable events beyond political stock markets. Traders increased prices correlated with things they wanted to happen in the laboratory markets (Forsythe et al., 1999). Beside prices, biases also affected the distribution of holdings across traders.

Oliven and Rietz (2004) also provide evidence of irrational trader behavior in a political stock market which cannot be explained by traders' biases. Traders are buying and selling at prices that are not the best available and are violating arbitrage restrictions. Despite this irrational behavior, markets are found to be remarkably efficient. Oliven and Rietz (2004) find dramatic differences in mistake rates between market-making and price-taking traders. Market-making traders are far less mistake-prone and in the end determine prices. This explains why markets can be efficient despite irrational trader behavior. Market makers profit from the other traders' mistakes and thus have an incentive to set efficient prices. Such erroneous actions, however, are not discussed in the following.

6.2. Description of the Data

This section describes the data which is used to study the correlation between the traders' nationality and their trading behavior. First of all, the analysis focuses on the STOCCER championship market which was already described in detail in Section 3.2.1. Contracts of all 32 national soccer teams were traded in this market. It was chosen for the following analysis because it was the most liquid market and the only market which was running continuously for several weeks. Furthermore, it is well suited to study the influence of the traders' nationality on their shareholdings and trading behavior since contracts of all national teams were traded in this market and the payoff of contracts depended on the overall performance of the teams. If biases related to the traders' country of origin existed, they should thus be observed in this market.

Every action of traders was recorded in the STOCCER championship market. Full information about the trading activity, i.e. orders and trades, and traders' shareholdings is available or can be calculated for any point in time. Moreover, the traders' nationality is known since they provided information about their country of origin during the registration process. As described in Section 3.2.4, traders originated from 72 different countries around the world. Countries with a substantial number of traders were Germany, Switzerland, USA, Belgium, Austria, UK, China, and Italy. The number of traders from other countries is too small to allow for a meaningful analysis of traders' biases. Out of the eight aforementioned countries, the following analysis is restricted to countries which were taking part in the FIFA World Cup 2006. Hence, data on shareholdings as well as trading activity is analyzed to study biases of traders coming from Germany, Switzerland, USA, UK, and Italy.

Beyond the STOCCER championship market, trading prices of two other prediction markets which also traded contracts on national teams were collected during the FIFA World Cup 2006. The first market, Ballstreet, was operated in Germany and focused on German traders. Since the web pages of Ballstreet were not translated it is quite unlikely that traders from other countries were joining this market. The second market, TradeSports, is targeted at traders coming from the US and UK. Trading prices from these two markets and STOCCER are compared to examine whether prices differ across prediction markets if the predominant majority of traders originates from different countries.

6.3. Results

Traders in STOCCER are expected to be overly optimistic about their national team's likely success and to interpret news with respect to their national team more favorably than other traders. Thus, they should overestimate the likely success of their national team and make larger investments (number of contracts held) in their national team. Section 6.3.1 describes how the traders' nationality affected their shareholdings in the STOCCER championship market. Subsequently, Section 6.3.2 shows the influence of the trader's nationality on their trading behavior. Prediction markets which focus on traders coming from one country can also be expected to exhibit a bias in favor of the corresponding country. Section 6.3.3 therefore contrasts trading prices of three prediction markets with different target groups of traders.

6.3.1. Traders' Nationality and Shareholdings

Similar to investors in financial markets who commonly allocate a large fraction of their portfolio to domestic investments, traders in the STOCCER championship market should hold more contracts of their country's national soccer team if they overestimate its likely success. Table 16 shows the average number of contracts held by traders originating from Germany, Switzerland, USA, UK, and Italy in the corresponding national teams at the market close on July 9th 2006[66]. Swiss traders, for instance, hold an average of about 1,153 contracts of the Swiss national team. They hold fewer contracts in the other four countries. On average across all 32 contracts traded in the market, Swiss traders hold only about 471 contracts.

Table 16: Traders' nationality and shareholdings in teams (July 9th 2006)

		AVERAGE NUMBER OF CONTRACTS					
		Germany	**Switzerland**	**USA**	**UK**	**Italy**	**Average**
TRADERS' NATIONALITY	**Germany**	401.97	214.02	326.26	323.84	324.75	311.74
	Switzerland	189.39	1153.06	592.93	262.11	396.83	471.30
	USA	218.86	95.39	387.39	377.06	268.29	213.18
	UK	70.00	73.33	60.00	1347.60	543.87	446.30
	Italy	79.69	114.54	226.08	79.92	1406.54	277.71

As a matter of fact, traders from all of these countries on average hold more shares in their own national team than in any of the other five teams. They also hold more contracts of their national team compared to the average team out of the 32 national soccer teams participating in the FIFA World Cup.

Figure 26 further highlights this bias by contrasting the average number of contracts held in the team of the traders' home country with the average number of contracts held across all teams on July 9th 2006[67]. It can be seen that traders from Germany,

[66] The biases which are observed here are not specific for this point in time. Table 27 in Appendix C, for example, contrasts the traders' nationality and shareholdings on June 23rd 2006, i.e. before the final rounds of the tournament started.
[67] Figure 31 in Appendix C shows the same comparison with data from June 23rd 2006, i.e. before the final rounds of the tournament started.

Switzerland, USA, UK, and Italy indeed hold more contracts of their national team than of other teams. On average, the 1,306 traders coming from these five countries held about 546 contracts of their own national team compared to 336 contracts across all 32 teams[68]. The difference between the number of contracts held by traders in their national team and the number of contracts held across all teams is significant (Mann-Whitney U test, p-value < 0.001)[69].

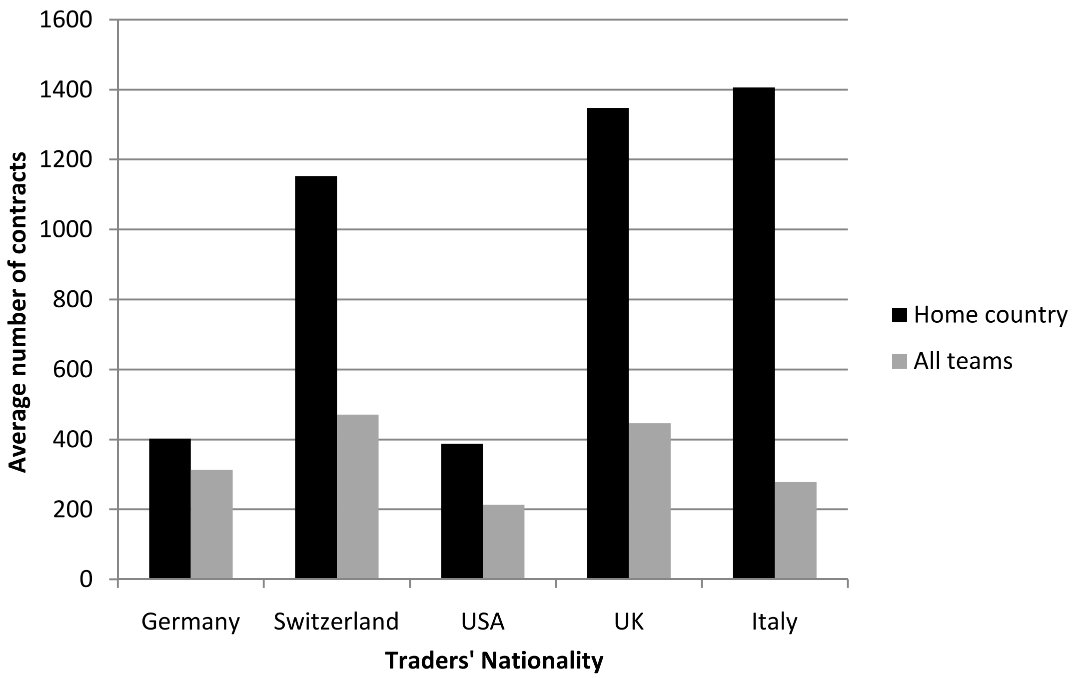

Figure 26: Shareholdings in home country and across all teams (July 9th 2006)

As a consequence, traders were biased in terms of holding more contracts of their own national soccer team than of other teams in the STOCCER championship markets. This can presumably be attributed to traders overestimating the likely success of their national team. If traders are more optimistic about their team than other traders, they should be willing to buy contracts at higher prices and thus also hold more contracts of their team than other traders.

[68] The standard deviation is 1503.72 for the contracts of the home country and 797.44 for all contracts.
[69] For more information on the Mann-Whitney U test see Mann and Whitney (1947).

6.3.2. Traders' Nationality and Trading Behavior

Biases observed in the traders' shareholdings result from their trading behavior. This section therefore studies how biases resulting from the traders' nationality impact their trading behavior in the STOCCER championship market. Since traders hold more contracts of their own national team there should be a larger proportion of net buyers among traders coming from the corresponding country compared to the proportion of net buyers among traders coming from other countries.

Table 17 shows the number and proportion of traders who purchased the contracts of the soccer teams from Germany, Switzerland, USA, UK, and Italy. For each contract, the traders are split up into two groups. The first group of traders comprises all traders coming from the country corresponding to the respective contract while the second group comprises all remaining traders. To give an example, there were 540 German traders who traded the contract "Germany". 413 out of these 540 traders bought at least one contract, i.e. the 127 remaining active traders only sold the contract. The proportion of German traders who bought the contract "Germany" thus is about 76 per cent whereas only about 57 per cent of non-German traders bought contracts of the German national team.

Table 17: Traders' nationality and proportion of buyers

Contracts	Traders' nationality	# active traders	#traders who purchased	% of traders who purchased	p-value[70]
Germany	Germany	540	413	76.48%	<0.001
	Other	188	107	56.91%	
Switzerland	Switzerland	122	112	91.80%	<0.001
	Other	471	243	51.59%	
USA	USA	16	12	75.00%	0.006
	Other	591	245	41.46%	
UK	UK	9	6	66.67%	0.584
	Other	646	482	74.61%	
Italy	Italy	7	7	100.00%	0.102
	Other	619	448	72.37%	

For four out of five contracts under investigation, the proportion of traders who purchased a contract was higher among traders coming from the corresponding country

[70] Chi-square test for difference in proportion of traders who purchased the corresponding contract

compared to the remaining traders. Merely in case of the UK, the proportion of traders who purchased is a little higher among non-UK traders than among UK traders. The difference in the proportion of traders is statistically significant for the contracts of Germany, Switzerland, and the United States of America (Pearson's chi-square test, see last column of Table 17)[71]. However, for the two contracts with a very small number of traders coming from the corresponding countries, i.e. UK and Italy, this difference is not statistically significant.

Table 18 follows the same idea but now shows the number and proportion of traders who sold the contracts of the five soccer teams. Again, the traders per contract are split up into the same two groups. For all five contracts, the proportion of traders who sold a contract was lower among traders coming from the corresponding country compared to the remaining traders. The difference in the proportion of traders is once more statistically significant for the contracts of Germany, Switzerland, and the United States of America (Pearson's chi-square test, see last column of Table 18). However, for the two contracts with a very small number of traders coming from the corresponding countries, i.e. UK and Italy, this difference is also not statistically significant.

Table 18: Traders' nationality and proportion of sellers

Contracts	Traders' nationality	# active traders	# traders who sold	% of traders who sold	p-value[72]
Germany	Germany	540	343	63.52%	0.001
	Other	188	132	70.21%	
Switzerland	Switzerland	122	58	47.54%	<0.001
	Other	471	385	81.74%	
USA	USA	16	8	50.00%	<0.001
	Other	591	530	89.68%	
UK	UK	9	4	44.44%	0.207
	Other	646	417	64.55%	
Italy	Italy	7	3	42.86%	0.170
	Other	619	416	67.21%	

Overall, the traders' nationality seems to influence the proportion of traders who are buying and selling contracts. The proportion of traders buying a contract at all is larger among traders coming from the corresponding country compared to other traders and,

[71] For more information on Pearson's chi-square test see e.g. Cowan (1998)
[72] Chi-square test for difference in proportion of traders who sold the corresponding contract

vice versa, the proportion of traders selling a contract is lower among traders coming from the corresponding country compared to other traders.

Yet, the number of net buyers among the two groups of traders is even more worthy of note than the number of traders who are buying and selling contracts at all. Table 19 therefore compares the proportion of traders with net purchases among traders coming from the corresponding country to the proportion of traders with net purchases from other countries for each of the five contracts.

Table 19: Traders' nationality and proportion of traders with net purchases

Contracts	Traders' nationality	# active traders	# traders with net purchases	% of traders with net purchases	p-value[73]
Germany	Germany	540	301	55.74%	<0.001
	Other	188	83	44.15%	
Switzerland	Switzerland	122	93	76.23%	<0.001
	Other	471	148	31.42%	
USA	USA	16	10	62.50%	<0.001
	Other	591	127	21.49%	
UK	UK	9	6	66.67%	0.345
	Other	646	329	50.93%	
Italy	Italy	7	5	71.43%	0.308
	Other	619	323	52.18%	

As can be seen in Table 19, there is indeed a larger proportion of net buyers among traders coming from the corresponding country compared to the proportion of net buyers among traders coming from other countries for all the contracts under investigation. The difference in the proportion of traders with net purchases is once more statistically significant for the contracts of Germany, Switzerland, and the United States of America (Pearson's chi-square test, see last column of Table 19). For the two contracts UK and Italy with a very small number of traders coming from the corresponding countries the difference is again not statistically significant.

All in all, the traders' nationality influences their trading behavior. The differences in the proportion of net buyers can most likely be attributed to traders overestimating the likely success of their national team. They are more optimistic about their team than

[73] Chi-square test for difference in proportion of traders with net purchases

other traders and thus are more likely to become net buyers of contracts related to their national soccer team.

6.3.3. Target Groups and Trading Prices

Beyond biases observed within one prediction market, one can also expect to find differences in trading prices between different prediction markets which focus on specific target groups, i.e. traders coming from different countries in case of the FIFA World Cup 2006. Trading prices should be biased in favor of a soccer team if the predominant majority of traders originate from the same country. In the following, trading prices from Ballstreet, which is targeted at German traders, and TradeSports, which is targeted at traders coming from the US and UK, are compared to trading prices of STOCCER in order to examine how prices differ across these prediction markets.

Figure 27 depicts the development of the rank at which the German soccer team is seen according to trading prices in STOCCER, Ballstreet, and TradeSports over time. Ballstreet traders, for instance, expect Germany to become the World Champion most of the time while Germany is expected to perform much worse by TradeSports traders. In the end, Germany did not win the World Cup but made the third place.

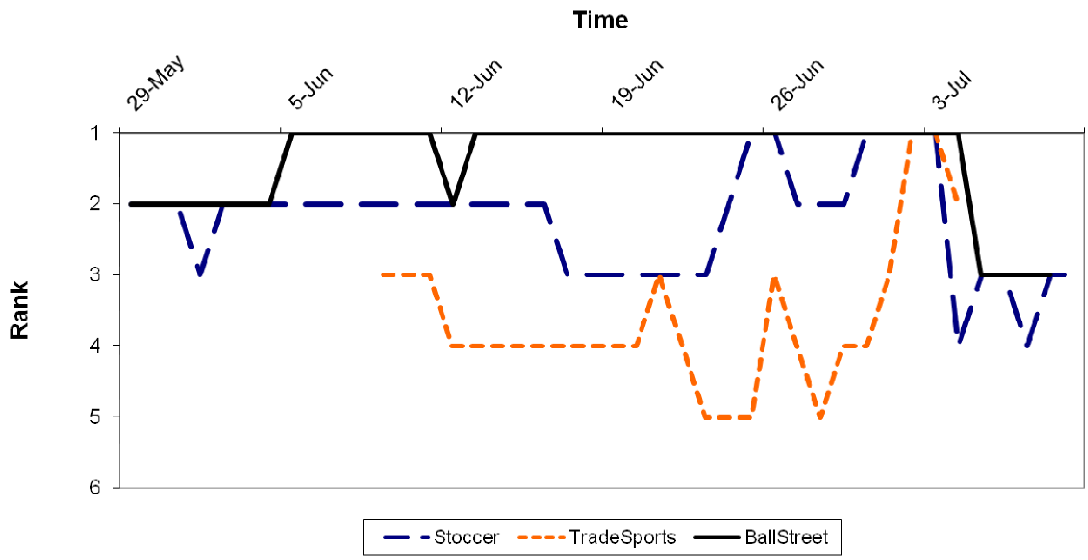

Figure 27: Trading prices of contract "Germany"

Traders' biases could explain these differences in expectations. In case of Ballstreet, almost all of the traders are supposed to be Germans and they are thus relatively optimistic about the success of the German soccer team. Trading in TradeSports, in contrast, is not dominated by Germans and consequently there is less optimism with regard to the German soccer team. In STOCCER, a large proportion of traders were Germans but other traders may have partly corrected the bias towards Germany. This would explain why the rank expected according to trading prices in STOCCER is between Ballstreet and TradeSports.

Similar to the bias towards Germany which is observed in Ballstreet, Figure 28 shows a bias towards the UK soccer team in case of TradeSports where trading is most likely dominated by British and American traders. The development of the rank at which the UK soccer team is seen according to trading prices over time shows less optimism towards the UK team in markets which are dominated by German traders. In this case, however, STOCCER most of the time ranks the UK team even lower than Ballstreet although there are German traders only in Ballstreet and at least few traders from the UK in the STOCCER championship market.

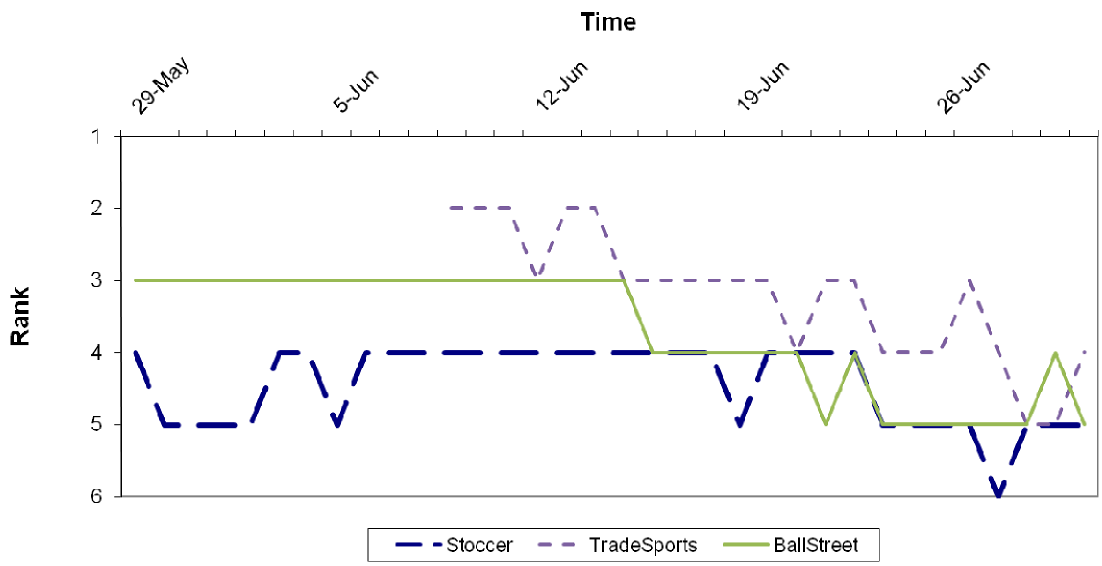

Figure 28: Trading prices of contract "UK"

If biases were observed beyond the German and the UK team in the above-mentioned prediction markets, they could rather not be explained by the nationality of the dominating group of traders. However, such noticeable biases cannot be found for other countries. One example is given in Figure 29 which depicts the development of the rank at which the French soccer team was seen according to trading prices in STOCCER, Ballstreet, and TradeSports during the World Cup. Ranks do not differ that much from one prediction market to another market. A similarly close correspondence between predictions based on trading prices of STOCCER, Ballstreet, and TradeSports can be found for almost all other contracts which are not expected to exhibit any biases.

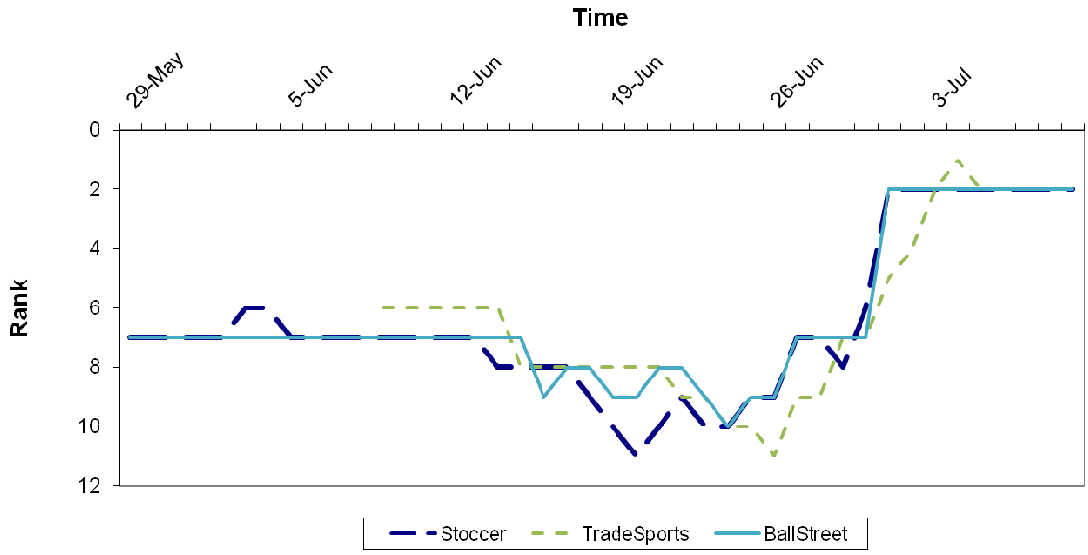

Figure 29: Trading prices of contract "France"

Differences in trading prices between prediction markets which focus on specific target groups are consequently indeed found here. A study based on a larger number of contracts traded, though, would be required to examine such differences between prediction markets in more depth.

6.4. Discussion of Results

The results reported in this chapter provide evidence that traders were biased in the STOCCER championship market. The traders' nationality influenced their trading behavior. Traders held more contracts of their own national soccer team than traders of a different nationality. Furthermore, the proportion of net buyers for all the contracts under investigation was found to be larger among traders coming from the corresponding country compared to the proportion of net buyers among traders coming from other countries.

These results are in line with earlier findings in the field of political stock markets. Forsythe et al. (1992) found that traders are buying and selling contracts of US presidential candidates in a manner which is correlated with their preferences, i.e. supporters of a candidate buy more contracts of this candidate than they sell. Forsythe et al. (1992) attributed the observed biases to failures in the traders' information-processing capabilities. However, attempts of manipulation could also have explained the traders' behavior in political stock markets. The results reported in this chapter contribute to the literature by demonstrating that such biases can also be found in field of application where traders cannot influence the outcome. In case of STOCCER, traders cannot influence the outcome of soccer matches or the performance of their national soccer team. Thus, manipulation cannot serve as an explanation for the traders' behavior in the STOCCER championship market. Failures in the traders' information-processing capabilities for that reason can in fact be seen as a plausible explanation for the trading behavior which was found in STOCCER.

Interestingly, the predictions of the STOCCER championship market were found to be very accurate (cp. Chapter 4) despite the biases which were found when looking at traders individually. Presumably, biases of a group of traders such as the traders coming from a certain country can be compensated by the remaining traders as long as the proportion of traders with biases in favor of the same contract is not too large. Similar to this, Hanson et al. (2006) found that subjects in an experimental market compensated for the bias in offers from manipulators who were submitting higher price offers by setting a different threshold at which they were willing to accept trades. As a result, the distortionary effects of manipulation were cancelled out in the experiment.

This also has important implications for selecting traders of prediction markets. Traders' biases most likely do not distort prediction accuracy if other traders are compensating for these biases. Prediction market operators thus have to ensure that not all traders exhibit the same bias. Otherwise, traders' biases could indeed distort trading prices and thereby also the prediction accuracy. The results reported in Section 6.3.3 suggest that the proportion of traders with the same bias was rather large in Ballstreet and TradeSports.

6.5. Summary

This chapter provided evidence of traders' biases in prediction markets. Earlier research on the home bias in financial markets and biases in political stock markets suggests that traders in prediction markets should exhibit substantial biases. Data from the STOCCER championship market is employed to study the influence of the traders' nationality on their shareholdings and trading behavior. Moreover, trading prices from Ballstreet and TradeSports were collected and compared to STOCCER prices to examine whether prices differ across prediction markets if the predominant majority of traders originate from different countries.

The results showed that the traders' nationality indeed influenced their trading behavior. Traders held more contracts of their own national soccer team than traders of a different nationality. Furthermore, the proportion of net buyers for all the contracts under investigation was found to be larger among traders coming from the corresponding country compared to the proportion of net buyers among traders coming from other countries. In addition, trading prices of prediction markets which focus on specific target groups seem to differ due to differences in the traders' biases.

The chapter concluded with a discussion of the results and their implications for inviting and selecting traders of prediction markets to avoid that the traders' biases distort prediction accuracy.

7. Conclusions and Future Work

This chapter concludes the work at hand by firstly summarizing the work in Section 7.1. Thereby, the research questions, the main contributions, and the implications of the results are reviewed. Then, Section 7.2 briefly outlines directions of future work which is closely related to the research questions addressed in the work at hand.

7.1. Summary of Contributions and Review of Work

The main objective of this work was to demonstrate the predictive power of markets in general and in the field of sports forecasting in particular. Moreover, the research which was presented concerning traders' biases and incentive schemes is valuable for designing future prediction markets.

The present work therefore makes the following contributions:

1. It provides evidence of markets' prediction accuracy and thereby contributes to the literature with the first empirical comparison of play-money prediction markets and predictions based on historic data or betting odds in the field of sports forecasting.
2. It analyzes the impact of different incentive schemes on the accuracy of prediction markets. In a field experiment, a rank-order tournament outperforms the fixed payment as well as the performance-compatible incentive scheme in terms of prediction accuracy.
3. It provides evidence of traders' biases in prediction markets. In a sports prediction market, the traders' nationality was found to influence their shareholdings as well as their trading behavior.

The work has proceeded in several steps to present these contributions. Chapter 1 motivated the work, raised the research questions which were addressed, and presented the structure of the work as well as related publications.

Chapter 2 gave a definition of prediction markets and explained their operational principle as well as their theoretical foundations. It also discussed the key design elements of prediction markets which have to be considered by market engineers. Moreover, Chapter 2 presented current fields of application of prediction markets.

Chapter 3 then described a 2006 FIFA World Cup prediction market called STOCCER. Most of the data which was used to answer the research questions raised in the work comes from STOCCER. For this reason the FIFA World Cup 2006 itself, the contracts that were traded, the trading mechanisms, the incentive schemes, the group of traders, as well as the software platform were described in detail.

Chapter 4 examined the accuracy of prediction markets in general and in the field of sports forecasting in particular, more precisely for predicting the outcomes of soccer matches during the FIFA World Cup 2006. The results showed that play-money prediction markets outperformed a random predictor and forecasts that are based on historic data about the success of national soccer teams. Moreover, prediction markets were found to be on a level with betting odds from professional bookmakers which are known to be very accurate. Beyond the comparison of prediction accuracy, Chapter 4 also studied whether pure arbitrage opportunities existed in these markets and whether traders tried to exploit illiquidity by taking on the role of market makers in prediction markets.

Chapter 5 studied the impact of different incentive schemes on prediction accuracy. It elaborated on the question whether or not prediction markets with performance-related incentives perform better than markets with flat payments and how these performance-related incentives should be designed. The highest correlation between the relative frequency of outcome and trading prices was found in case of a rank-order tournament where the most successful traders were paid depending on their ordinal rank in a group of traders.

Chapter 6 analyzed how the traders' nationality influenced their shareholdings and their trading behavior in a FIFA World Cup 2006 prediction market. The results suggested that there was a correlation between the traders' nationality and the number of contracts they held of different national teams. Moreover, Chapter 6 showed that the proportion of net buyers for all the contracts under investigation was larger among traders coming from the corresponding country compared to the proportion of net buyers among traders coming from other countries. In spite of these results predictions from these markets were surprisingly accurate.

Based on the results reported in this work, the three research questions which were posed in the introduction can be answered briefly as follows.

(I) How well do markets predict the future?

Earlier empirical research on prediction markets is used to substantiate their predictive power in several fields of application. In this work, data collected from prediction markets for the FIFA World Cup 2006 demonstrates their predictive power in the field of sports forecasting. The results show that the play-money prediction market STOCCER for the FIFA World Cup 2006 was about as accurate as betting odds. Betting odds, in turn, are known to be very accurate predictors. Moreover, the markets clearly outperformed predictions based on the FIFA world ranking as well as the random predictions. Overall, prediction markets can thus be considered an extremely accurate forecasting method.

(II) How to design incentive schemes for play-money prediction markets?

Well-designed incentive schemes are needed to encourage participation and information revelation in play-money prediction markets. In this work, three widely-used incentive schemes were compared with regard to their impact on the accuracy of predictions in a field experiment. The results show that rank-order tournaments are a suitable incentive scheme in case of risk-averse traders. The competitive environment in the corresponding treatment overrides risk aversion and in so doing leads to the best results in terms of prediction accuracy. Due to the risk aversion average trading prices were by far too low in case of the performance-compatible incentive schemes. The fixed payment scheme was also found not to be too well-suited since there was no trading activity for several events and also no significant correlation between trading prices and outcome frequencies.

(III) How do traders' biases impact their trading behavior?

Prediction markets aggregate and reveal the information traders have. Individuals, however, exhibit substantial information processing or judgment biases. This work studied how the traders' nationality influenced their shareholdings and their trading behavior in a FIFA World Cup 2006 prediction market. The results show that the traders' nationality in fact influenced their trading behavior. Traders held more contracts of their own national soccer team than traders of a different nationality. Furthermore, the proportion of net buyers for all the contracts under investigation was found to be larger among traders coming from the corresponding country compared to the proportion of net buyers among traders coming from other countries. Furthermore, trading prices of prediction markets which focus on specific target groups seem to differ due to differences in the traders' biases.

All in all, this work provides further evidence of the predictive power of markets. The markets work extremely well despite traders' biases. This once more substantiates the impressive robustness of prediction markets. Moreover, this work emphasizes the role of incentive schemes for the successful operation of prediction markets.

7.2. Future Work

Prediction markets have continuously gained importance in academia and industry over the last couple of years. Nevertheless, it is a rather new field of research and numerous open questions still need to be tackled. Several interesting streams of research are closely related to the work at hand.

Fields of application

The work at hand provided evidence of markets' prediction accuracy in the field of sports forecasting. So far, most of the research comparing the accuracy of prediction markets to other forecasting methods focused on field of application where information is dispersed among a large group of traders. Thus, it is interesting to extend this stream of research to other fields of application where relevant information is only available to a limited number of experts and to study how well prediction markets work under such circumstances. This would also allow for examining whether adding uninformed traders to a market with few well-informed experts distorts trading prices and thus harms prediction accuracy.

Incentive schemes

Concerning incentive schemes of play-money prediction markets, the explanation which was given in this work for the poor performance of the performance-compatible payment was the traders' risk aversion. In order to study the impact of risk-aversion on the performance of different incentive schemes, one could conduct a similar field experiments as the one described in this work with risk-seeking traders. Moreover, it would be extremely relevant for transferring the results into practice to study the traders' risk attitude in prediction markets where traders are self-selected. If traders are similarly risk averse as in the field experiment, the results from this work should be taken into consideration when designing incentive schemes for play-money markets.

Traders' biases

This work provided evidence of traders' biases in a prediction market for the FIFA World Cup 2006. First of all, future research is required to investigate whether similar biases can also be found in fields of application where biases are less obvious compared to sports and political elections. Furthermore, field and lab experiments have to be conducted to study when and to which extent traders' biases influence trading prices. For example, one has to find out what proportion of non-biased traders or how many traders with other biases are required in order to correct for the majority's biases. Thereby, it does not matter whether biases result from political preferences, nationality, manipulation, or the traders' nescience.

Combining prediction markets with established forecasting methods

The track record of prediction markets suggests that markets may help to better foresee future developments and trends. Yet, other forecasting methods should not always be replaced by prediction markets. Markets can rather be thought of as a supplement to existing forecasting methods since they can be seen as a tool for continuous monitoring of developments. Moreover, prediction markets are useful to motivate creative thinking and idea generation as well as to identify knowledgeable traders which can afterwards be recruited as experts for alternative forecasting methods such as the Delphi technique.

Prediction markets can also be combined with voting mechanisms or crowd-based innovation. The "Open Innovation Markets" developed by Xpree[74], for example, make use of the *wisdom of crowds* to facilitate crowd sourcing, crowd ranking, and crowd analysis of innovations. The idea is to brainstorm as a community, vote on the ideas to

[74] http://www.xpree.com

rank them, and then forecast key metrics using a prediction market. Such combinations of several forecasting methods should be considered when aiming at improving a company's foresight capabilities.

Appendix A

Table 20: Betting odds from wetten.de

Match (Team 1 – Team 2)	Odds 1	0	2	Result (Team 1 – Team 2)
Germany - Costa Rica	1.26	5.45	13.00	4-2
Poland - Ecuador	1.90	3.35	4.40	0-2
England - Paraguay	1.62	3.55	6.45	1-0
Trinidad & Tobago - Sweden	14.00	5.45	1.25	0-0
Argentina - Ivory Coast	1.55	3.75	7.05	2-1
Serbia & Montenegro - Netherlands	4.45	3.30	1.90	0-1
Mexico - Iran	1.55	3.80	6.85	3-1
Angola - Portugal	10.00	4.70	1.35	0-1
Australia - Japan	2.60	3.20	2.80	3-1
USA - Czech Republic	4.45	3.30	1.90	0-3
Italy - Ghana	1.58	3.50	7.35	2-0
South Korea - Togo	2.00	3.25	4.05	2-1
France - Switzerland	1.70	3.35	6.00	0-0
Brazil - Croatia	1.40	4.50	8.50	1-0
Spain - Ukraine	1.85	3.30	4.75	4-0
Tunisia - Saudi Arabia	1.83	3.35	4.80	2-2
Germany - Poland	1.55	3.85	6.70	1-0
Ecuador - Costa Rica	1.82	3.55	4.50	3-0
England - Trinidad & Tobago	1.20	6.50	15.00	2-0
Sweden - Paraguay	1.85	3.40	4.55	1-0
Argentina - Serbia & Montenegro	1.55	3.50	6.50	6-0
Netherlands - Ivory Coast	1.80	3.50	4.70	2-1
Mexico - Angola	1.45	4.35	7.45	0-0
Portugal - Iran	1.35	4.50	8.50	2-0
Czech Republic - Ghana	1.60	3.65	5.75	0-2
Italy - USA	1.45	3.80	7.80	1-1
Japan - Croatia	5.60	3.60	1.60	0-0
Brazil - Australia	1.25	5.20	11.00	2-0
France - South Korea	1.40	4.10	8.00	1-1
Togo - Switzerland	8.75	4.00	1.30	0-2
Saudi Arabia - Ukraine	8.75	4.35	1.35	0-4
Spain - Tunisia	1.30	4.50	10.00	3-1
Ecuador - Germany	7.50	4.25	1.40	0-3
Costa Rica - Poland	4.00	3.30	1.75	1-2
Sweden - England	3.65	2.30	2.55	2-2

Match (Team 1 – Team 2)	Odds 1	Odds 0	Odds 2	Result (Team 1 – Team 2)
Paraguay - Trinidad & Tobago	1.90	3.55	3.30	2-0
Portugal - Mexico	2.40	2.50	3.50	2-1
Iran - Angola	2.55	3.30	2.45	1-1
Netherlands - Argentina	3.40	3.10	2.05	0-0
Ivory Coast - Serbia & Montenegro	1.95	3.40	3.40	3-2
Czech Republic - Italy	3.55	2.80	2.15	0-2
Ghana - USA	2.20	3.30	2.90	2-1
Japan - Brazil	11.25	5.50	1.20	1-4
Croatia - Australia	2.10	3.30	3.10	2-2
Saudi Arabia - Spain	12.00	5.50	1.18	0-1
Ukraine - Tunisia	1.60	3.60	5.00	1-0
Togo - France	11.00	5.00	1.20	0-2
Switzerland - South Korea	1.90	2.95	3.85	2-0
Germany – Sweden	1.60	3.45	5.50	2-0
Argentina – Mexico	1.40	4.00	8.00	1-1
England – Ecuador	1.50	3.60	7.00	1-0
Portugal – Netherlands	3.00	3.05	2.35	1-0
Italy – Australia	1.45	3.75	7.50	1-0
Switzerland – Ukraine	2.40	3.00	2.90	0-0
Brazil – Ghana	1.25	5.15	10.00	3-0
Spain – France	2.35	3.05	3.00	1-3
Germany – Argentina	2.60	3.10	2.60	1-1
England – Portugal	2.15	3.10	3.35	0-0
Italy – Ukraine	1.55	3.45	6.35	3-0
Brazil – France	1.75	3.20	4.75	0-1
Germany – Italy	2.20	3.00	3.30	0-0
Portugal – France	3.65	3.05	2.05	0-1
Germany – Portugal	1.75	3.40	4.40	3-1
Italy – France	2.50	2.80	3.00	1-1

Table 21: Betting odds from ODDSET

Match (Team 1 – Team 2)	Odds 1	0	2	Result (Team 1 – Team 2)
Germany - Costa Rica	1.20	4.00	9.00	4-2
Poland - Ecuador	1.75	2.85	3.40	0-2
England - Paraguay	1.45	2.90	5.45	1-0
Trinidad & Tobago - Sweden	9.00	4.00	1.20	0-0
Argentina - Ivory Coast	1.50	2.85	5.00	2-1
Serbia & Montenegro - Netherlands	3.60	2.85	1.70	0-1
Mexico - Iran	1.40	3.20	5.20	3-1
Angola - Portugal	7.50	3.50	1.25	0-1
Australia - Japan	1.80	2.90	3.15	3-1
USA - Czech Republic	3.45	2.80	1.75	0-3
Italy - Ghana	2.25	2.75	2.45	2-0
South Korea - Togo	1.30	3.40	6.50	2-1
France - Switzerland	1.55	2.85	4.50	0-0
Brazil - Croatia	1.40	3.10	5.50	1-0
Spain - Ukraine	1.75	2.80	3.50	4-0
Tunisia - Saudi Arabia	1.75	2.80	3.50	2-2
Germany - Poland	1.40	3.10	5.50	1-0
Ecuador - Costa Rica	1.80	2.80	3.30	3-0
England - Trinidad & Tobago	1.15	5.00	10.00	2-0
Sweden - Paraguay	1.75	2.85	3.40	1-0
Argentina - Serbia & Montenegro	1.50	3.00	4.60	6-0
Netherlands - Ivory Coast	1.65	2.80	4.00	2-1
Mexico - Angola	1.30	3.55	6.00	0-0
Portugal - Iran	1.30	3.55	6.00	2-0
Czech Republic - Ghana	1.35	3.25	6.00	0-2
Italy - USA	1.50	3.00	4.60	1-1
Japan - Croatia	1.20	4.00	8.25	0-0
Brazil - Australia	3.60	2.85	1.70	2-0
France - South Korea	1.35	3.25	6.00	1-1
Togo - Switzerland	5.00	3.30	1.40	0-2
Saudi Arabia - Ukraine	1.25	3.50	7.50	0-4
Spain - Tunisia	6.00	3.55	1.30	3-1
Ecuador - Germany	6.00	3.55	1.30	0-3
Costa Rica - Poland	4.00	3.10	1.55	1-2
Sweden - England	3.00	2.35	2.20	2-2
Paraguay - Trinidad & Tobago	1.70	3.25	3.10	2-0
Portugal - Mexico	3.00	2.85	1.90	2-1
Iran - Angola	1.85	2.90	3.00	1-1
Netherlands - Argentina	2.10	2.40	3.10	0-0

Match (Team 1 – Team 2)	**Odds**			**Result (Team 1 – Team 2)**
	1	**0**	**2**	
Ivory Coast - Serbia & Montenegro	2.40	2.90	2.20	3-2
Czech Republic - Italy	3.25	2.60	1.90	0-2
Ghana - USA	2.00	2.80	2.80	2-1
Japan - Brazil	7.50	4.20	1.20	1-4
Croatia - Australia	2.00	2.75	2.85	2-2
Saudi Arabia - Spain	10.00	4.25	1.15	0-1
Ukraine - Tunisia	1.90	2.60	3.25	1-0
Togo - France	10.00	4.25	1.15	0-2
Switzerland - South Korea	1.50	3.00	4.60	2-0
Germany – Sweden	1.60	3.00	4.00	2-0
Argentina – Mexico	1.35	3.25	6.00	1-1
England – Ecuador	1.35	3.25	6.00	1-0
Portugal – Netherlands	2.70	2.80	2.05	1-0
Italy – Australia	1.40	3.00	6.00	1-0
Switzerland – Ukraine	2.20	2.80	2.50	0-0
Brazil – Ghana	1.20	4.00	8.25	3-0
Spain – France	2.15	2.75	2.60	1-3
Germany – Argentina	2.35	2.75	2.35	1-1
England – Portugal	1.95	2.75	3.00	0-0
Italy – Ukraine	1.45	3.00	5.10	3-0
Brazil – France	1.60	2.85	4.15	0-1
Germany – Italy	1.95	2.75	3.00	0-0
Portugal – France	3.15	2.70	1.90	0-1
Germany – Portugal	1.65	2.90	3.75	3-1
Italy – France	2.30	2.60	2.60	1-1

Table 22: Positions of competing teams in the FIFA ranking (May 2006)

Match (Team 1 – Team 2)	Rank Team 1	Team 2	Result (Team 1 – Team 2)
Germany - Costa Rica	19	26	4-2
Poland - Ecuador	29	39	0-2
England - Paraguay	10	33	1-0
Trinidad & Tobago - Sweden	47	16	0-0
Argentina - Ivory Coast	9	32	2-1
Serbia & Montenegro - Netherlands	47	3	0-1
Mexico - Iran	4	23	3-1
Angola - Portugal	57	7	0-1
Australia - Japan	42	18	3-1
USA - Czech Republic	5	2	0-3
Italy - Ghana	13	48	2-0
South Korea - Togo	29	61	2-1
France - Switzerland	8	35	0-0
Brazil - Croatia	1	23	1-0
Spain - Ukraine	5	45	4-0
Tunisia - Saudi Arabia	21	34	2-2
Germany - Poland	19	29	1-0
Ecuador - Costa Rica	39	26	3-0
England - Trinidad & Tobago	10	47	2-0
Sweden - Paraguay	16	33	1-0
Argentina - Serbia & Montenegro	9	47	6-0
Netherlands - Ivory Coast	3	32	2-1
Mexico - Angola	4	57	0-0
Portugal - Iran	7	23	2-0
Czech Republic - Ghana	2	48	0-2
Italy - USA	13	5	1-1
Japan - Croatia	18	23	0-0
Brazil - Australia	1	42	2-0
France - South Korea	8	29	1-1
Togo - Switzerland	61	35	0-2
Saudi Arabia - Ukraine	34	45	0-4
Spain - Tunisia	5	21	3-1
Ecuador - Germany	39	19	0-3
Costa Rica - Poland	26	29	1-2
Sweden - England	16	10	2-2
Paraguay - Trinidad & Tobago	33	47	2-0
Portugal - Mexico	7	4	2-1
Iran - Angola	23	57	1-1
Netherlands - Argentina	3	9	0-0

Match (Team 1 – Team 2)	**Rank Team 1**	**Rank Team 2**	**Result (Team 1 – Team 2)**
Ivory Coast - Serbia & Montenegro	32	47	3-2
Czech Republic - Italy	2	13	0-2
Ghana - USA	48	5	2-1
Japan - Brazil	18	1	1-4
Croatia - Australia	23	42	2-2
Saudi Arabia - Spain	34	5	0-1
Ukraine - Tunisia	45	21	1-0
Togo - France	61	8	0-2
Switzerland - South Korea	35	29	2-0
Germany – Sweden	19	16	2-0
Argentina – Mexico	9	4	1-1
England – Ecuador	10	39	1-0
Portugal – Netherlands	7	3	1-0
Italy – Australia	13	42	1-0
Switzerland – Ukraine	35	45	0-0
Brazil – Ghana	1	48	3-0
Spain – France	5	8	1-3
Germany – Argentina	19	9	1-1
England – Portugal	10	7	0-0
Italy – Ukraine	13	45	3-0
Brazil – France	1	8	0-1
Germany – Italy	19	13	0-0
Portugal – France	7	8	0-1
Germany – Portugal	19	7	3-1
Italy – France	13	8	1-1

Table 23: Trading activity of market makers relative to all traders

Contract	#MM	#MM-TX / #TX (%)	MM-TradVol / TradVol (%)
Angola	45	76.19%	89.51%
Argentina	59	83.34%	82.42%
Australia	54	77.70%	77.33%
Brazil	56	84.26%	87.41%
Costa Rica	45	76.28%	91.46%
Cote d'Ivoire	41	79.21%	87.57%
Croatia	47	83.54%	89.96%
Czech Republic	39	82.04%	86.63%
Ecuador	42	82.66%	87.61%
England	53	85.83%	85.77%
France	77	83.74%	81.98%
Germany	81	81.74%	80.43%
Ghana	50	80.01%	78.31%
Iran	25	76.61%	83.00%
Italy	59	84.62%	83.38%
Japan	32	78.92%	81.28%
Korea Republic	47	81.59%	87.14%
Saudi Arabia	36	79.48%	86.24%
Mexico	50	82.88%	82.12%
Netherlands	51	86.73%	89.22%
Paraguay	36	80.21%	90.10%
Poland	37	79.68%	88.66%
Portugal	49	85.25%	81.73%
Serbia & Montenegro	32	80.16%	90.84%
Spain	59	84.20%	82.56%
Sweden	45	84.98%	87.79%
Switzerland	46	83.03%	85.54%
Togo	32	78.87%	88.60%
Trinidad & Tobago	43	77.54%	81.92%
Tunisia	36	82.02%	94.56%
Ukraine	54	82.24%	82.12%
USA	44	80.55%	82.04%

#MM: Number of market makers

#TX: Number of trades

TradVol: Trading volume

#MM-TX: Number of trades by market makers

MM-TradVol: Trading volume of market makers

Table 24: Number of market makers and trading activity per contract

Contract	# MM	# TX	Trading Volume
Angola	45	2822	2906207.80
Argentina	59	3397	16518302.03
Australia	54	2628	5669446.43
Brazil	56	3456	21245499.70
Costa Rica	45	2188	1768325.72
Cote d'Ivoire	41	2491	3101242.95
Croatia	47	2284	4051174.70
Czech Republic	39	2311	5415731.57
Ecuador	42	2538	5698810.33
England	53	2633	10684352.88
France	77	3524	19028177.09
Germany	81	3494	19461286.03
Ghana	50	2756	6698774.88
Iran	25	2129	1911784.25
Italy	59	2809	15022296.44
Japan	32	2182	2658963.66
Korea Republic	47	2173	3822122.80
Saudi Arabia	36	2071	1588805.83
Mexico	50	2576	7509094.91
Netherlands	51	2404	7744212.78
Paraguay	36	1971	2717072.52
Poland	37	2224	3173347.09
Portugal	49	2658	13111409.97
Serbia & Montenegro	32	2142	2919919.26
Spain	59	2772	11381556.92
Sweden	45	2150	5552289.44
Switzerland	46	2151	5149225.96
Togo	32	2087	1550324.84
Trinidad & Tobago	43	2297	2770702.86
Tunisia	36	2124	3124018.13
Ukraine	54	2528	7253846.15
USA	44	2432	4209720.01

#MM: Number of market makers

#TX: Number of trades

Appendix B

Instruction 1: Instructions sent to subjects via e-mail when trading started (incentive scheme: rank-order tournament)

Hallo «Vorname» «Nachname»,

es ist soweit: Die Handelsplattform für unser Experiment steht Ihnen unter der Adresse http://exp.iw.uni-karlsruhe.de/ zur Verfügung. Hier können Sie virtuelle Aktien handeln und echtes Geld verdienen! Wie Sie entlohnt werden finden Sie nachfolgend unter „Entlohnung der Experimentteilnehmer".

Ausgangssituation
Zu Beginn haben Sie einen Geldbestand von «Erstausstattung_Geld» Stoccer-Euro und besitzen Aktien im Wert von «Erstausstattung_Turnieraktien» Stoccer-Euro. Während des Experiments werden Sie weitere Aktien im Wert von «Erstausstattung_Spielaktien» Stoccer-Euro erhalten. Wir werden Sie in jedem Fall noch genauer per Email informieren, wenn sich neue Aktien in Ihrem Depot befinden.

Anmeldung
Der Handel findet ausschließlich auf unserer speziellen Handelsplattform statt, die unter der folgenden Adresse erreicht werden kann: http://exp.iw.uni-karlsruhe.de/

Bitte loggen Sie sich mit Ihrem Benutzernamen und Passwort ein:
Benutzername: «benutzername»
Passwort: «passwort»

Marktauswahl
Direkt nach der Anmeldung können Sie den Markt auswählen, in dem Sie handeln wollen. Anfangs ist nur der Markt „Fußball-Weltmeister" aktiv. Nach und nach werden wir weitere Märkte zu den einzelnen Spielen starten.

Auszahlungsregel
Die Aktien der 32 WM-Mannschaften werden bis zum 9. Juli 2006 gehandelt. Nach dem Finale der Fußball-WM 2006 werden die Aktien abhängig von Turniererfolg der Mannschaften zu folgenden Preisen (in STOCCER-Euro) bewertet:

Turniererfolg	*Wert*
Weltmeister	50
Vize-Weltmeister	30
Ausscheiden im Halbfinale	20
Ausscheiden im Viertelfinale	10
Ausscheiden im Achtelfinale	5

So wird die Aktie des Weltmeisters z.B. nach Ende der WM 50 STOCCER-Euro wert sein, die Aktie einer Mannschaft, welche die Vorrunde übersteht, dann aber im Achtelfinale verliert, 5 STOCCER-Euro.

Handel

Handeln können Sie unter dem Menüpunkt "Aktien handeln". Dort finden Sie neben den Orderbüchern mit je drei aktuellen Geboten anderer Händler auch Ihren aktuellen Depotbestand und die Maske zur Erteilung von Aufträgen zum Kauf- und Verkauf von Anteilen. Zu Beginn handeln Sie mit Aktien der 32 WM-Mannschaften. Die Preise Ihrer Kauf- und Verkaufaufträge sind abhängig von Ihren Erwartungen. Liegt bspw. der aktuelle Marktpreis von Brasilien bei 27 und Sie gehen davon aus, dass Brasilien mindestens das Finale erreichen wird, so würden Sie gemäß obiger Auszahlungsregel sofort Aktien von Brasilien kaufen, da Brasilien Ihrer Meinung nach mindestens Vize-Weltmeister wird.
Zur eigentlichen Durchführung eines Kaufs oder Verkaufs wählen Sie eine Mannschaft und tragen die gewünschte Stückzahl und den Preis ein. Der Auftrag wird dann zu dem von Ihnen gewählten Preis oder einem für Sie besseren Preis ausgeführt, sobald es eine entsprechende Order auf der Marktgegenseite gibt. Selbstverständlich kann Ihr Auftrag nur ausgeführt werden, wenn Sie genügend Aktien zum Verkauf beziehungsweise genügend Geld zum Ankauf haben. Zusätzlich können Sie die Gültigkeitsdauer Ihrer Order bestimmen. Eventuell auftretende Fehler bei der Auftragserteilung werden im Log-Bereich des Handelsbildschirms angezeigt. Eine ausführliche Anleitung finden Sie auf der Webseite unter http://www.stoccer.de/index.php?id=36.

Entlohnung der Experimentteilnehmer

Sie handeln in einer Gruppe mit insgesamt 20 Teilnehmern. Am Ende des Experiments wird auf Basis des Depotwerts eine Rangfolge bestimmt und es werden folgende Geldbeträge ausgezahlt:
500€ für den ersten Platz (höchster Depotwert), 300€ für den zweiten Platz und 200€ für den dritten Platz unter allen 20 Händlern.

Bitte beachten Sie, dass ein Mindesttransaktionsvolumen von «Mindesttransaktionsvolumen» Stoccer-Euro pro Woche besteht. Sollten Sie weniger als dieses Volumen (Anzahl * Preis) handeln, wird am Ende des Experiments eine Gebühr von 15,00€ für jede Woche (jeweils Freitag bis Donnerstag) erhoben, in der das Mindesttransaktionsvolumen nicht erreicht wurde. Die Gebühr wird von Ihrer Experimentauszahlung abgezogen – die Auszahlung kann natürlich nicht negativ werden. Bitte beachten Sie, dass sich das Transaktionsvolumen aus allen ausgeführten Transaktionen unterschiedlicher Händler errechnet und Geschäfte, in denen Sie mit sich selbst handeln, nicht berücksichtigt werden.

Transaktionen und offene Aufträge

Im Menüpunkt "Aufträge und Transaktionen" finden Sie sowohl bereits ausgeführte als auch noch offene Aufträge. Die noch offenen Aufträge können Sie hier jederzeit löschen.

Handel von Portfolios
Alternativ zum oben beschriebenen Handel einzelner Anteile mit anderen Marktteilnehmern können Sie unter dem Menüpunkt "Portfolios handeln" sog. Basisportfolios kaufen und verkaufen. Im Markt "Fußball-Weltmeister" beinhaltet ein Portfolio eine Aktie von jeder Mannschaft. Im Dialog geben Sie bitte an, ob Sie kaufen oder verkaufen möchten und tragen die Anzahl an Portfolios ein, die Sie handeln möchten. Wenn Sie genügend Aktien zum Verkauf beziehungsweise genügend Geld zum Ankauf haben, wird der Auftrag ausgeführt.
Der Preis für ein Portfolio ist marktabhängig und liegt bei 200 STOCCER-Euro. Damit entspricht der Preis exakt der Summe der Auszahlungen der einzelnen Mannschaftsaktien. In den Märkten zu einzelnen Spielen wird der Portfoliopreis bei 100 STOCCER-Euro liegen.

Berechnung des Depotwerts
Ihre Platzierung innerhalb Ihrer Gruppe ist abhängig vom Depotwert. Dieser ergibt sich aus der Summe Ihres Geldbestands und des aktuellen Werts aller von Ihnen gehaltenen Aktien. Nach Marktschluss werden die Aktien entsprechend der Auszahlungsregel (siehe "Auszahlungsregel) bewertet.

Sollten Sie Fragen zum Ablauf des Experiments haben, zögern Sie nicht, die Experimentleitung unter Matthias.Burghardt@iism.uni-karlsruhe.de oder Stefan.Luckner@iism.uni-karlsruhe.de zu kontaktieren.

Vielen Dank für Ihre Teilnahme und viel Erfolg beim Handeln!

Matthias Burghardt
Stefan Luckner

IISM Universität Karlsruhe (TH)
www.iw.uni-karlsruhe.de

Instruction 2: Instructions sent to subjects when the first match markets were launched

Hallo «Vorname» «Nachname»,

neben dem Markt „Fußball-Weltmeister“ gibt es im Experiment weitere Märkte zu den einzelnen Spielen. In jedem dieser „Spielmärkte“ werden drei verschiedene Aktien gehandelt:

1. Team 1 gewinnt
2. Team 2 gewinnt
3. Unentschieden

Die Auszahlung für die Aktie des eintretenden Ereignisses beträgt **100 STOCCER-Euro**, für alle anderen Aktien 0 STOCCER-Euro. Demnach kostet ein Portfolio 100 STOCCER-Euro. Erwarten Sie also bspw., dass Brasilien gegen Ghana mit 90%iger Wahrscheinlichkeit gewinnen wird, so kaufen Sie Brasilien Aktien bis zu einem Preis von 90 STOCCER-Euro und verkaufen Sie zu Preise über 90 STOCCER-Euro.
Die Märkte enden direkt nach dem Ende der 2. Halbzeit. Die Basis zur Bewertung ist der **Spielstand nach Ende der 2. Halbzeit** (d.h. 90 min. Spielzeit + Nachspielzeit), also **ohne Verlängerung/Elfmeterschießen**!

---
1. Beispiel: Beim Spiel Deutschland gegen Schweden gewinnt Deutschland regulär mit dem Abpfiff der 2. Halbzeit mit 2:0. Dann ist die Aktie "Deutschland gewinnt" 100 STOCCER-Euro wert, die Aktien "Schweden gewinnt" und "Unentschieden" sind 0 STOCCER-Euro wert.
2. Beispiel: Es steht es 1:1 nach Ende der 2. Halbzeit, d.h. das Spiel geht in die Verlängerung. Da die Aktien nach Ende der 2. Halbzeit bewertet werden, ist die Auszahlung für die Aktie "Unentschieden" 100 STOCCER-Euro, für die Aktien "Deutschland gewinnt" und "Schweden gewinnt" je 0 STOCCER-Euro.

Matthias Burghardt
Stefan Luckner

IISM Universität Karlsruhe (TH)
www.iw.uni-karlsruhe.de

Table 25: Distribution of trading prices in the three treatments

Trading price	FP	RO	DV
0-20	0.32	0.19	0.52
20-40	0.27	0.29	0.22
40-60	0.20	0.17	0.22
60-80	0.07	0.21	0.04
80-100	0.14	0.14	0.00

FP: Fixed payment

RO: Rank-order tournament

DV: Deposit value

Table 26: Relative frequencies of outcome of contracts

FP

Price range	Relative outcome frequency
0-20	42,86%
20-40	8,33%
40-60	11,11%
60-80	66,67%
80-100	66,67%

RO

Price range	Relative outcome frequency
0-20	25,00%
20-40	41,67%
40-60	28,57%
60-80	55,56%
80-100	66,67%

DV

Price range	Relative outcome frequency
0-20	16,67%
20-40	40,00%
40-60	30,00%
60-80	0,00%
80-100	0,00%

Figure 30: Lottery-choice decisions in the field experiment

Entscheidung	**Option A** Würfelwurf ergibt	Option zahlt	Ich wähle Option A	**Option B** Würfelwurf ergibt	Option zahlt	Ich wähle Option B
1	1 2,3,4,5,6,7,8,9 oder 10	100,00 Euro 80,00 Euro		1 2,3,4,5,6,7,8,9 oder 10	192,50 Euro 5,00 Euro	
2	1 oder 2 3,4,5,6,7,8,9 oder 10	100,00 Euro 80,00 Euro		1 oder 2 3,4,5,6,7,8,9 oder 10	192,50 Euro 5,00 Euro	
3	1,2 oder 3 4,5,6,7,8,9 oder 10	100,00 Euro 80,00 Euro		1,2 oder 3 4,5,6,7,8,9 oder 10	192,50 Euro 5,00 Euro	
4	1,2,3 oder 4 5,6,7,8,9 oder 10	100,00 Euro 80,00 Euro		1,2,3 oder 4 5,6,7,8,9 oder 10	192,50 Euro 5,00 Euro	
5	1,2,3,4 oder 5 6,7,8,9 oder 10	100,00 Euro 80,00 Euro		1,2,3,4 oder 5 6,7,8,9 oder 10	192,50 Euro 5,00 Euro	
6	1,2,3,4,5 oder 6 7,8,9 oder 10	100,00 Euro 80,00 Euro		1,2,3,4,5 oder 6 7,8,9 oder 10	192,50 Euro 5,00 Euro	
7	1,2,3,4,5,6 oder 7 8,9 oder 10	100,00 Euro 80,00 Euro		1,2,3,4,5,6 oder 7 8,9 oder 10	192,50 Euro 5,00 Euro	
8	1,2,3,4,5,6,7 oder 8 9 oder 10	100,00 Euro 80,00 Euro		1,2,3,4,5,6,7 oder 8 9 oder 10	192,50 Euro 5,00 Euro	
9	1,2,3,4,5,6,7,8 oder 9 10	100,00 Euro 80,00 Euro		1,2,3,4,5,6,7,8 oder 9 10	192,50 Euro 5,00 Euro	
10	1,2,3,4,5,6,7,8,9 oder 10 ---	100,00 Euro 80,00 Euro		1,2,3,4,5,6,7,8,9 oder 10 ---	192,50 Euro 5,00 Euro	

Appendix C

Table 27: Traders' nationality and shareholdings in teams (June 23rd 2006)

		AVERAGE NUMBER OF CONTRACTS					
		Germany	**Switzerland**	**USA**	**UK**	**Italy**	**Average**
TRADERS' NATIONALITY	**Germany**	405.24	244.73	391.85	333.58	330.94	330.10
	Switzerland	180.16	1027.52	294.40	270.05	372.06	442.67
	USA	219.56	97.30	395.14	350.08	307.26	235.18
	UK	75.00	71.43	64.29	946.86	372.50	446.76
	Italy	87.36	78.27	260.09	87.36	1900.55	299.22

Figure 31: Shareholdings in home country and across all teams (June 23rd 2006)

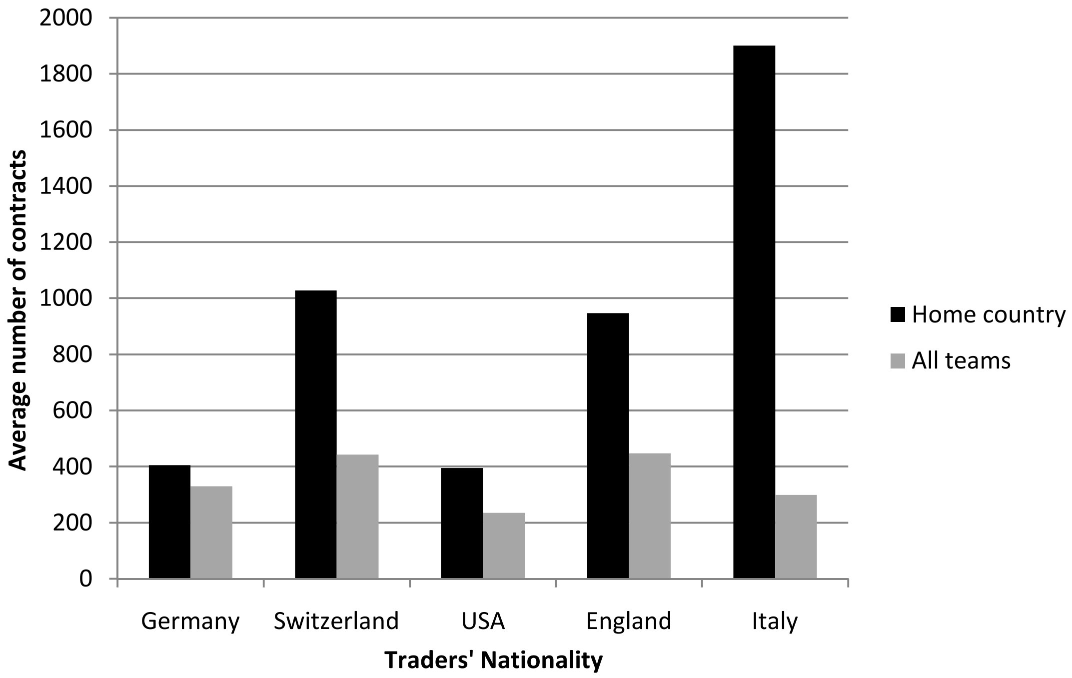

Bibliography

AHEARNE, A. G., GRIEVER, W. L. & WARNOCK, F. E. (2004) Information costs and home bias: an analysis of US holdings of foreign equities. *Journal of International Economics,* 62 (2), 313-336.

ALI, M. M. (1977) Probability and Utility Estimates for Racetrack Bettors. *Journal of Political Economy,* 85 (4), 803-815.

ANTWEILER, W. & ROSS, T. (1998) The 1997 UBC Election Stock Market. *Canadian Business Economics,* 6 (2), 15-22.

BARTELS, L. (1987) *Presidential Primaries and The Dynamics of Public Choice,* Princeton, NJ, Princeton University Press.

BARTLETT, M. S. (1937) Properties of Sufficiency and Statistical Tests. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences,* 160 (901), 268-282.

BECKMANN, K. & WERDING, M. (1996) "Passauer Wahlbörse": Information Processing in a Political Market Experiment. *Kyklos,* 49 (2), 171-204.

BERG, J. E., FORSYTHE, R., NELSON, F. & RIETZ, T. A. (2001) Results from a Dozen Years of Election Futures Markets Research. IN PLOTT, C. & SMITH, V. L. (Eds.) *Handbook of Experimental Economic Results.* Amsterdam, Elservier Science.

BERG, J. E., FORSYTHE, R. & RIETZ, T. A. (1996) What Makes Markets Predict Well? Evidence from the Iowa Electronic Markets. IN ALBERS, W., GÜTH, W., HAMMERSTEIN, P., MOLDOVANU, B. & VAN DAMME, E. (Eds.) *Understanding Strategic Interaction: Essays in Honor of Reinhard Selten.* New York, Springer.

BERG, J. E., FORSYTHE, R. & RIETZ, T. A. (1997) The Iowa Electronic Market. IN PAXSON, D. & WOOD, D. (Eds.) *Blackwell Encyclopedic Dictionary of Finance.* Oxford, UK, Blackwell.

BERG, J. E., NELSON, F. & RIETZ, T. A. (2003) Accuracy and Forecast Standard Error of Prediction Markets. Working Paper, College of Business Administration, University of Iowa.

BERG, J. E. & RIETZ, T. A. (2003) Prediction Markets as Decision Support Systems. *Information Systems Frontiers,* 5 (1), 79-93.

BERG, J. E. & RIETZ, T. A. (2006) The Iowa Electronic Markets: Stylized Facts and Open Issues. IN HAHN, R. & TETLOCK, P. C. (Eds.) *Information Markets: A*

New Way of Making Decisions in the Public and Private Sectors. Washington D.C., AEI Press.

BEWLEY, T. F. (1995) A Depressed Labor Market as Explained by Participants. *American Economic Review,* 85 (2), 250-254.

BOHM, P. & SONNEGARD, J. (1999) Political Stock Markets and Unreliable Polls. *The Scandinavian Journal of Economics,* 101 (2), 205-222.

BONDARENKO, O. & BOSSAERTS, P. (2000) Expectations and learning in Iowa. *Journal of Banking and Finance,* 24 (9), 1535-1555.

BONNER, S. E., HASTIE, R., SPRINKLE, G. B. & YOUNG, S. M. (2000) A Review of the Effects of Financial Incentives on Performance in Laboratory Tasks: Implications for Management Accounting. *Journal of Management Accounting Research,* 12 (1), 19-64.

BRÜGGELAMBERT, G. (2004) Information and efficiency in political stock markets: using computerized markets to predict election results. *Applied Economics,* 36 (7), 753-768.

BULL, C., SCHOTTER, A. & WEIGELT, K. (1987) Tournament and Piece Rates: An Experimental Study. *Journal of Political Economy,* 95 (1), 1-33.

CAIN, M., LAW, D. & PEEL, D. (2000) The Favourite-Longshot Bias and Market Efficiency in UK Football Betting. *Scottish Journal of Political Economy,* 47 (1), 25-36.

CAMERER, C. F. & HOGARTH, R. M. (1999) The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework. *Journal of Risk and Uncertainty,* 19 (1-3), 7-42.

CAMERON, J., BANKO, K. M. & PIERCE, W. D. (2001) Pervasive Negative Effects of Rewards on Intrinsic Motivation: The Myth Continues. *The Behavior Analyst,* 24 (1), 1-44.

CHAN, N., DAHAN, E., KIM, A., LO, A. & POGGIO, T. (2002) Securities Trading of Concepts (STOC). Working Paper, Massachusetts Institute of Technology.

CHEN, K.-Y. & PLOTT, C. R. (2002) Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem. *Social Science Working Paper No.1131.* Pasadena, California Institute of Technology.

CHEN, Y., CHU, C.-H., MULLEN, T. & PENNOCK, D. M. (2005) Information markets vs. opinion pools: an empirical comparison. *Proceedings of the 6th ACM Conference on Electronic commerce.* Vancouver, BC, Canada.

CHRISTIANSEN, J. D. (2007) Prediction Markets: Practical Experiments in Small Markets and Behaviours Observed. *Journal of Prediction Markets,* 1 (1), 17-41.

CLAR, G. (2003) Forecasting Options for the Future - to Gain Foresight to Select and Shape Them. *Journal of Forecasting,* 22 (2-3), 83-91.

COOPER, I. & KAPLANIS, E. (1994) Home Bias in Equity Portfolios, Inflation Hedging, and International Capital Market Equilibrium. *Review of Financial Studies,* 7 (1), 45-60.

COVAL, J. D. & MOSKOWITZ, T. J. (1999) Home Bias at Home: Local Equity Preference in Domestic Portfolios. *Journal of Finance,* 54 (6), 2045-2073.

COWAN, G. (1998) *Statistical Data Analysis*, Oxford University Press.

CUHLS, K. (2003) From Forecasting to Foresight Processes - New Participative Foresight Activities in Germany. *Journal of Forecasting,* 22 (2-3), 93-111.

DEBNATH, S., PENNOCK, D. M., GILES, C. L. & LAWRENCE, S. (2003) Information incorporation in online in-Game sports betting markets. *Proceedings of the 4th ACM conference on Electronic commerce.* San Diego, CA, USA, ACM Press.

ERIKSON, R. S. & WLEZIEN, C. (2006) Are Political Markets Really Superior to Polls as Election Predictors? *Annual Meeting of the Midwest Political Science Association.* Chicago.

FAMA, E. F. (1970a) Efficient Capital Markets: A Review of Theory and Empirical Work. *Journal of Finance,* 25, 383-417.

FAMA, E. F. (1970b) Efficient capital markets: A review of theory and empirical work. *Journal of Finance,* 25 (2), 383-417.

FAMA, E. F. (1991) Efficient Capital Markets: II. *Journal of Finance,* 46 (5), 1575-1617.

FIGLEWSKI, S. (1979) Subjective Information and Market Efficiency in a Betting Model. *Journal of Political Economy,* 87 (1), 75-88.

FILZMAIER, P., BEYRL, M., HAUSER, F. & HUBER, J. (2003) Wahlbörsen als interdisziplinäres Instrument der Sozialforschung. *SWS Rundschau,* 3, 387-410.

FORREST, D., GODDARD, J. & SIMMONS, R. (2005) Odds-setters as forecasters: The case of English football. *International Journal of Forecasting,* 21 (3), 551-564.

FORSYTHE, R., FRANK, M., KRISHNAMURTHY, V. & ROSS, T. W. (1995) Using Market Prices to Predict Election Results: The 1993 UBC Election Stock Market. *The Canadian Journal of Economics,* 28 (4), 770-793.

FORSYTHE, R., FRANK, M., KRISHNAMURTHY, V. & ROSS, T. W. (1998) Markets as Predictors of Election Outcomes: Campaign Events and Judgment Bias in the 1993 UBC Election Stock Market. *Canadian Public Policy,* 24 (3).

FORSYTHE, R., NELSON, F., NEUMANN, G. & WRIGHT, J. (1992) Anatomy of an Experimental Political Stock Market. *American Economic Review,* 82 (5), 1142-1161.

FORSYTHE, R., NELSON, F., NEUMANN, G. & WRIGHT, J. (1994) The 1992 Iowa Political Stock Market: September Forecasts. *The Political Methodologist,* 5 (2), 15-19.

FORSYTHE, R., RIETZ, T. A. & ROSS, T. (1999) Wishes, expectations and actions: a survey on price formation in election stock markets. *Journal of Economic Behavior and Organization,* 39 (1), 83-110.

FOWLER, J. H. (2006) Elections and Markets: The Effect of Partisanship, Policy Risk, and Electoral Margins on the Economy. *The Journal of Politics,* 68 (1), 89-103.

FRANKE, M., GEYER-SCHULZ, A. & HOSER, B. (2005) Analyzing Trading Behavior in Transaction Data of Electronic Election Markets. IN BAIER, D., DECKER, R. & SCHMIDT-THIEME, L. (Eds.) *Data Analysis and Decision Support.* Springer.

FRANKE, M., GEYER-SCHULZ, A. & HOSER, B. (2006) On the Analysis of Asymmetric Directed Communication Structures in Electronic Election Markets. IN BILLARI, F. C., FENT, T., PRSKAWETZ, A. & SCHEFFRAN, J. (Eds.) *Agent-Based Computational Modelling.* Physica.

FRENCH, K. R. & POTERBA, J. M. (1991) Investor Diversification And International Equity Markets. *American Economic Review,* 81 (2), 222-226.

GANDAR, J. M., DARE, W. H., BROWN, C. R. & ZUBER, R. A. (1998) Informed Traders and Price Variations in the Betting Market for Professional Basketball Games. *Journal of Finance,* 53 (1), 385-401.

GEYER-SCHULZ, A., LUCKNER, S., SCHRÖDER, J., SKIERA, B., SLAMKA, C. & WEINHARDT, C. (2007) Empirical Evaluation of Call Auctions in Prediction Markets. Working Paper, Universität Karlsruhe (TH).

GLASSMANN, D. A. & RIDDICK, L. A. (2001) What causes home asset bias and how should it be measured? *Journal of Empirical Finance,* 8 (1), 35-54.

GLOSTEN, L. R. & MILGROM, P. R. (1985) Bid, Ask and Transaction Prices in a Specialist Market With Heterogeneously Informed Traders. *Journal of Financial Economics,* 14, 71-100.

GNEEZY, U. & RUSTICHINI, A. (2000) Pay Enough Or Don'T Pay At All. *The Quarterly Journal of Economics,* 115 (3), 791-810.

GRUCA, T. S., BERG, J. E. & CIPRIANO, M. (2003) The Effect of Electronic Markets on Forecasts of New Product Success. *Information Systems Frontiers,* 5 (1), 95-105.

GÜRKAYNAK, R. S. & WOLFERS, J. (2006) Macroeconomic derivatives: an initial analysis of market-based macro forecasts, uncertainty, and risk. IN PISSARIDES, C. & FRANKEL, J. (Eds.) *NBER International Seminar on Macroeconomics.* MIT Press.

HANSEN, J., SCHMIDT, C. & STROBEL, M. (2004) Manipulation in Political Stock Markets - Preconditions and Evidence. *Applied Economics Letters,* 11 (7), 459-463.

HANSON, R. (1990a) Could Gambling Save Science? Encouraging an Honest Consensus. *Proceedings of the Eighth International Conference on Risk and Gambling.* London.

HANSON, R. (1990b) Market-Based Foresight: A Proposal. *Foresight Update,* 10, 1-4.

HANSON, R. (1992) Idea Futures: Encouraging an Honest Consensus. *Extropy,* 3 (2), 7-17.

HANSON, R. (1999) Decision Markets. *IEEE Intelligent Systems,* 14 (3), 16-19.

HANSON, R. (2003) Combinatorial Information Market Design. *Information Systems Frontiers,* 5 (1), 105-119.

HANSON, R. (2006) Foul Play in Information Markets. IN HAHN, R. & TETLOCK, P. C. (Eds.) *Information Markets: A New Way of Making Decisions in the Public and Private Sectors.* Washington D.C., AEI Press.

HANSON, R., OPREA, R. & PORTER, D. (2006) Information Aggregation and Manipulation in an Experimental Market. *Journal of Economic Behavior and Organization,* 60 (4), 449-459.

HARRISON, G. W., LAU, M. I. & RUTSTRÖM, E. E. (2007) Estimating Risk Attitudes in Denmark: A Field Experiment. *Scandinavian Journal of Economics,* 109 (2), 341-368.

HAUSCH, D. B., ZIEMBA, W. T. & RUBINSTEIN, M. (1981) Efficiency of the Market for Racetrack Betting. *Management Science,* 27 (12), 1435-1452.

HAYEK, F. (1945) The Use of Knowledge in Society. *American Economic Review,* 35 (4), 519-530.

HOLT, C. A. & LAURY, S. K. (2002) Risk Aversion and Incentive Effects. *American Economic Review,* 92 (5), 1644-1655.

HOLT, C. A. & LAURY, S. K. (2005) Risk Aversion and Incentive Effects: New Data Without Order Effects. *American Economic Review,* 95 (3), 902-912.

HOTELLING, H. & PABST, M. R. (1936) Rank Correlation and Tests of Significance Involving No Assumption of Normality. *The Annals of Mathematical Statistics,* 7 (1), 29-43.

HUBER, J. & HAUSER, F. (2005) Systematic mispricing in experimental markets – evidence from political stock markets. *Proceedings of the International Conference on Finance.* Kopenhagen, Denmark.

JACOBSEN, B., POTTERS, J., SCHRAM, A., VAN WINDEN, F. & WIT, J. (2000) (In)accuracy of a European political stock market: The influence of common value structures - Ambiguity and competence in choice under uncertainty. *European Economic Review,* 44 (2), 205-230.

JENSEN, M. C. (1978) Some Anomalous Evidence Regarding Market Efficiency. *Journal of Financial Economics,* 6 (2/3), 95-101.

KALAY, A., WEI, L. & WOHL, A. (2002) Continuous Trading or Call Auctions: Revealed Preferences of Investors at the Tel Aviv Stock Exchange. *Journal of Finance,* 57 (1), 523-542.

KANG, J.-K. & STULZ, R. M. (1997) Why is there a home bias? An analysis of foreign portfolio equity ownership in Japan. *Journal of Financial Economics,* 46 (1), 3-28.

KIVIAT, B. (2004) The End of Management? *Time,* 164 (2).

KOHN, A. (1993) Why Incentive Plans Cannot Work. *Hardvard Business Review,* 71 (5), 54-63.

KOU, S. G. & SOBEL, M. E. (2004) Forecasting the Vote: A Theoretical Comparison of Election Markets and Public Opinion Polls. *Political Analysis,* 12 (3), 277-295.

KRUSKAL, W. H. & WALLIS, W. A. (1952) Use of ranks in one-criterion variance analysis. *Journal of the American Statistical Association,* 47 (260), 583-621.

LAMARE, A. (2007) Hollywood Stock Exchange (HSX.com) Traders correctly picked 7 out of 8 Top Category Oscar Winners to continue its stellar record. Los Angeles, California, USA, Press Release, Hollywood Stock Exchange.

LAZEAR, E. P. & ROSEN, S. (1981) Rank-Order Tournaments as Optimum Labor Contracts. *The Journal of Political Economy,* 89 (5), 841-864.

LEDYARD, J. (2006) Designing Information Markets for Policy Analysis. IN HAHN, R. & TETLOCK, P. C. (Eds.) *Information Markets: A New Way of Making Decisions in the Public and Private Sectors.* Washington D.C., AEI Press.

LEIGH, A., WOLFERS, J. & ZITZEWITZ, E. (2007) Is There a Favorite-Longshot Bias in Election Prediction Markets? *Workshop on the Growth of Gambling and Prediction Markets.* Palm Desert, California.

LEVY, H. & SARNAT, M. (1970) International Diversification of Investment Portfolios. *American Economic Review,* 60 (4), 668-675.

LUCKNER, S. (2006a) Monetary Incentives in Prediction Markets. *European Regional Meeting of the Economic Science Association (ESA).* Nottingham, UK.

LUCKNER, S. (2006b) Prediction Markets: How Do Incentive Schemes Affect Prediction Accuracy? IN JENNINGS, N., KERSTEN, G., OCKENFELS, A. & WEINHARDT, C. (Eds.) *Dagstuhl Seminar Proceedings 06461 Negotiation and Market Engineering.* Schloss Dagstuhl, Germany.

LUCKNER, S. (2007a) A Field Experiment on Monetary Incentives in Prediction Markets. *2nd Workshop on Prediction Markets.* San Diego, California.

LUCKNER, S. (2007b) A Field Experiment on Monetary Incentives in Prediction Markets. *Doctoral Consortium, 8. Internationale Tagung Wirtschaftsinformatik 2007.* Karlsruhe, Germany.

LUCKNER, S. (2007c) Price Formation in Sports Prediction Markets – A Cross-Cultural Study. IN KERSTEN, G. E., RIOS, J. & CHEN, E. (Eds.) *Group Decision and Negotiation (GDN) 2007.* Montreal, Canada.

LUCKNER, S., SCHRÖDER, J. & SLAMKA, C. (2007) On the Forecast Accuracy of Sports Prediction Markets. IN GIMPEL, H., JENNINGS, N. R., KERSTEN, G., OCKENFELS, A. & WEINHARDT, C. (Eds.) *Negotiation and Market Engineering.* LNBIP 2, Springer.

LUCKNER, S. & WEINHARDT, C. (2007) How to Pay Traders in Information Markets? Results from a Field Experiment. *Journal of Prediction Markets,* 1 (2), 147-156.

MADHAVAN, A. (1992) Trading Mechanisms in Securities Markets. *Journal of Finance,* 47 (2), 607-641.

MALONEY, M. & MULHERIN, H. (2003) The Complexity of Price Discovery in an Efficient Market: The Stock Market Reaction to the Challenger Crash. *Journal of Corporate Finance,* 9 (4), 453-479.

MANGOLD, B., DOOLEY, M., FLAKE, G. W., HOFFMAN, H., KASTURI, T., PENNOCK, D. M. & DORNFEST, R. (2005) The Tech Buzz Game. *Computer,* 38 (7), 94-97.

MANN, H. B. & WHITNEY, D. R. (1947) On a test of whether one of two random variables is stochastically larger than the other. *Annals of Mathematical Statistics,* 18 (1), 50-60.

MCKELVEY, R. D. & PAGE, T. (1990) Public and Private Information: An Experimental Study of Information Pooling. *Econometrica,* 58 (6), 1321-1339.

MERTON, R. C. (1987) A Simple Model of Capital Market Equilibrium with Incomplete Information. *Journal of Finance,* 42 (3), 483-510.

NOHRIA, N. & STEWART, T. A. (2006) Risk, Uncertainty, and Doubt. *Harvard Business Review,* 84 (2), 39-40.

OLIVEN, K. & RIETZ, T. A. (2004) Suckers Are Born but Markets Are Made: Individual Rationality, Arbitrage, and Market Efficiency on an Electronic Futures Market. *Management Science,* 50 (3), 336-351.

ORTNER, G. (1997) Forecasting Markets - An Industrial Application: Part I. Working Paper, TU Vienna.

ORTNER, G. (2000) Aktienmärkte als Industrielles Vorhersagemodell. *Zeitschrift für Betriebswirtschaft (ZfB) - Ergänzungsheft,* 70 (1), 115-125.

ORTNER, G., STEPAN, A. & ZECHNER, J. (1995) Political Stock Markets - The Austrian Experiences. *Zeitschrift für Betriebswirtschaft (ZfB),* Ergänzungsheft 4/95, 123-135.

PÁSTOR, Ľ. (2000) Portfolio Selection and Asset Pricing Models. *Journal of Finance,* 55 (1), 179-223.

PENNOCK, D. M. (2004) A Dynamic Pari-Mutuel Market for Hedging, Wagering, and Information Aggregation. *ACM Conference on Electronic Commerce 2004.* New York, USA.

PENNOCK, D. M., LAWRENCE, S., GILES, C. L. & NIELSEN, F. A. (2000) The Power of Play: Efficiency and Forecast Accuracy of Web Market Games. *Technical Report 2000-168.* Princeton, NEC Research Institute.

PENNOCK, D. M., LAWRENCE, S., NIELSEN, F. A. & GILES, C. L. (2001a) Extracting collective probabilistic forecasts from web games. *Proceedings of the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.* San Francisco, CA, USA.

PENNOCK, D. M., LAWRENCE, S., NIELSEN, F. A. & GILES, C. L. (2001b) The real power of artificial markets. *Science,* 291 (5506), 987-988.

PLOTT, C. R. (2000) Markets as Information Gathering Tools. *Southern Economic Journal,* 67 (1), 1-15.

PLOTT, C. R. & SUNDER, S. (1982) Efficiency of experimental security markets with insider information: An application of rational expectations models. *Journal of Political Economy,* 90 (4), 663-698.

PLOTT, C. R. & SUNDER, S. (1988) Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets. *Econometrica,* 56 (5), 1085-1118.

POLGREEN, P. M., NELSON, F. D. & NEUMANN, G. R. (2007) Use of Prediction Markets to Forecast Infectious Disease Activity. *Clinical Infectious Diseases,* 44 (2), 272-279.

POPE, P. F. & PEEL, D. A. (1989) Information, Prices and Efficiency in a Fixed-Odds Betting Market. *Economica,* 56 (223), 323-341.

PRENDERGAST, C. (1999) The Provision of Incentives in Firms. *Journal of Economic Literature,* 37 (1), 7-63.

ROBINSON, S. (2001) Play Money? Not to these News Investors. *New York Times.* New York, 5 July.

ROLL, R. (1984) Orange Juice and Weather. *American Economic Review,* 74 (5), 861-880.

ROSENBLOOM, E. S. & NOTZ, W. W. (2006) Statistical Tests of Real-Money versus Play-Money Prediction Markets. *Electronic Markets - The International Journal,* 16 (1).

SALO, A. & CUHLS, K. (2003) Technology foresight—past and future. *Journal of Forecasting,* 22 (2-3), 79-82.

SCHMIDT, C. & WERWATZ, A. (2002) How well do markets predict the outcome of an event? The Euro 2000 soccer championships experiment. *Discussion Papers on Strategic Interaction.* Jena, Germany, Max Planck Institute for Research into Economic Systems.

SERVAN-SCHREIBER, E., WOLFERS, J., PENNOCK, D. & GALEBACH, B. (2004) Prediction Markets: Does Money Matter? *Electronic Markets - The International Journal,* 14 (3), 243-251.

SHILLER, R. J., KON-YA, F. & TSUTSUI, Y. (1991) Speculative behavior in the stock markets: evidence from the United States and Japan. NBER Working Paper 3613.

SKIERA, B. & SPANN, M. (2004) Opportunities of Virtual Stock Markets to Support New Product Development. IN ALBERS, S. (Ed.) *Cross-functional Innovation Management.* Wiesbaden, Gabler.

SMITH, M. A., PATON, D. & WILLIAMS, L. V. (2006) Market Efficiency in Person-to-Person Betting. *Economica,* 73 (292), 673-689.

SMITH, V. L. & WALKER, J. M. (1993) Monetary Rewards and Decision Cost in Experimental Economics. *Economic Inquiry,* 31 (2), 245-261.

SNOWBERG, E. & WOLFERS, J. (2007) Explaining the Favorite-Longshot Bias: Is it Risk-Love or Misperceptions. Working Paper, Stanford Graduate School of Business.

SPANN, M., ERNST, H., SKIERA, B. & SOLL, J. H. (2007) Identification of Lead Users for Consumer Products via Virtual Stock Markets. *Journal of Product Innovation Management (forthcoming).*

SPANN, M. & SKIERA, B. (2003) Internet-Based Virtual Stock Markets for Business Forecasting. *Management Science,* 49 (10), 1310-1326.

SUNSTEIN, C. R. (2006) Deliberation and Information Markets. IN HAHN, W. & TETLOCK, P. C. (Eds.) *Information Markets: A New Way of Making Decisions in the Public and Private Sectors.* Washington D.C., AEI Press.

SUROWIECKI, J. (2004) *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*, Doubleday, USA.

TESAR, L. L. & WERNER, I. M. (1995) Home bias and high turnover. *Journal of International Money and Finance,* 14 (4), 467-492.

THALER, R. & ZIEMBA, W. (1988) Parimutuel Betting Markets: Racetracks and Lotteries. *Journal of Economic Perspectives,* 2 (2), 161-174.

TZIRALIS, G. & TATSIOPOULOS, I. (2007a) Prediction Markets: An Extended Literature Review. *Journal of Prediction Markets,* 1 (1), 75-91.

TZIRALIS, G. & TATSIOPOULOS, I. (2007b) Prediction markets: an information aggregation perspective to the forecasting problem. *World Review of Entrepreneurship, Management and Sustainable Development,* 3 (3/4), 251-259.

URBAN, G. L. & HAUSER, J. R. (1993) *Design and Marketing of New Products*, Prentice-Hall.

VAN BRUGGEN, G. H., SPANN, M., LILIEN, G. L. & SKIERA, B. (2006) Institutional Forecasting: The Performance of Thin Virtual Stock Markets. *ERIM Report Series Reference No. ERS-2006-028-MKT.* Rotterdam, Erasmus Research Institute of Management (ERIM).

VLASTAKIS, N., DOTSIS, G. & MARKELOS, R. N. (2006) Beating the Odds: Arbitrage and Winning Strategies in the European Football Betting Market. *European Financial Management Association Annual Meeting.* Madrid, Spain.

WEINHARDT, C., HOLTMANN, C. & NEUMANN, D. (2003) Market Engineering. *Wirtschaftsinformatik,* 45 (6), 635-640.

WEINHARDT, C., NEUMANN, D. & HOLTMANN, C. (2006a) Computer-aided Market Engineering. *Communications of the ACM,* 49 (7), 79.

WEINHARDT, C., VAN DINTHER, C., GRUNENBERG, M., KOLITZ, K., KUNZELMANN, M., MÄKIÖ, J., WEBER, I. & WELTZIEN, H. (2006b) *CAME-Toolsuite meet2trade - auf dem Weg zum Computer Aided Market Engineering,* Karlsruhe, Universitätsverlag Karlsruhe.

WEINHARDT, C., VAN DINTHER, C., KOLITZ, K., MÄKIÖ, J. & WEBER, I. (2005) meet2trade: A generic electronic trading platform. *4th Workshop on e-Business (WEB 2005).* Las Vegas, USA.

WOLFERS, J. & ZITZEWITZ, E. (2004) Prediction Markets. *Journal of Economic Perspectives,* 18 (2), 107-126.

WOLFERS, J. & ZITZEWITZ, E. (2005) Five Open Questions About Prediction Markets. IN HAHN, R. & TETLOCK, P. C. (Eds.) *Information Markets: A New Way of Making Decisions in the Public and Private Sectors.* Washington D.C., AEI Presss.

WOLFERS, J. & ZITZEWITZ, E. (2006a) Interpreting Prediction Market Prices as Probabilities. NBER Working Paper No. 12200.

WOLFERS, J. & ZITZEWITZ, E. (2006b) Prediction Markets in Theory and Practice. NBER Working Paper No. 12083.

WOODLAND, L. M. & WOODLAND, B. M. (1994) Market Efficiency and the Favorite-Longshot Bias: The Baseball Betting Market. *Journal of Finance,* 49 (1), 269-279.