

IZA DP No. 6720

Prediction Markets for Economic Forecasting

Erik Snowberg
Justin Wolfers
Eric Zitzewitz

July 2012

Prediction Markets for Economic Forecasting

Erik Snowberg

Caltech and NBER

Justin Wolfers

*Wharton, CEPR,
CESifo, IZA and NBER*

Eric Zitzewitz

Dartmouth and NBER

Discussion Paper No. 6720

July 2012

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Prediction Markets for Economic Forecasting^{*}

Prediction markets – markets used to forecast future events – have been used to accurately forecast the outcome of political contests, sporting events, and, occasionally, economic outcomes. This chapter summarizes the latest research on prediction markets in order to further their utilization by economic forecasters. We show that prediction markets have a number of attractive features: they quickly incorporate new information, are largely efficient, and impervious to manipulation. Moreover, markets generally exhibit lower statistical errors than professional forecasters and polls. Finally, we show how markets can be used to both uncover the economic model behind forecasts, as well as test existing economic models.

JEL Classification: C5, G14

Keywords: prediction markets, forecasting

Corresponding author:

Justin Wolfers
Business and Public Policy Dept.
The Wharton School
University of Pennsylvania
1456 Steinberg Hall-Dietrich Hall
3620 Locust Walk
Philadelphia, PA 19104-6372
USA
E-mail: jwolfers@wharton.upenn.edu

^{*} This was prepared for the Handbook of Economic Forecasting, Volume 2. We dedicate this chapter to the memory of John Delaney, founder and CEO of Intrade.com. An innovator and source of inspiration for many prediction market researchers, he will be sorely missed.

1 Introduction

Market prices, in the form of gambling odds, have been used to forecast events since at least the beginning of the sixteenth century. The use of such prices had a heyday in the early twentieth century, when gambling odds on elections were printed daily in newspapers such as *The New York Times*. This was followed by a decline in popularity, due largely to the advent of scientific polling (Rhode and Strumpf, 2004, 2008). Scientific interest in market prices as tools for forecasting was kindled in the second half of the twentieth century by the efficient markets hypothesis and experimental economics (Plott and Sunder, 1982, 1988; Berg et al., 2008). This scientific foundation, coupled with advances in telecommunications—which allowed prices to be shared in real time across companies and the globe—has led to resurgent interest in using markets for forecasting (Snowberg, Wolfers and Zitzewitz, 2007b).¹

Despite this long history, and markets' proven track-record of providing accurate forecasts of uncertain events, prediction markets—markets used to forecast future events—are largely unused in economic forecasting. There are some exceptions: the difference in yields between inflation protected and standard government bonds is used to forecast inflation, and futures contracts are sometimes used to forecast commodity prices, such as the price of oil. However, as other chapters in this volume reveal, these uses are accompanied by concerns about what else, other than information about future events, may be lurking in market prices (Alquist and Vigfusson, 2012; Duffee, 2012; Wright and Faust, 2012; Zhou and Rapach, 2012).

This chapter brings together the latest research on prediction markets to further their utilization by economic forecasters. We begin by providing a description of standard types of prediction markets, and an heuristic framework useful in understanding why prediction markets do, and sometimes do not, make useful predictions. We then show that, in practice, prediction markets often have a number of attractive features: they quickly incorporate new information, are largely efficient, and impervious to manipulation. Moreover, markets

¹See Wolfers and Zitzewitz (2008a) for an overview, and Tziralis and Tatsiopoulos (2007) for an exhaustive literature review.

generally outperform professional forecasters and polls. Finally, we argue that examining co-movement in market prices, through, for example, event studies, can be used to shed light on the underlying economic model of market participants. We conclude with a short list of open questions that may be of particular interest to economic forecasters.

2 Types of Prediction Markets

A prediction market is generally implemented as a wager (or contract) that pays off if a particular outcome, such as an economic indicator taking a particular value y , occurs. Assuming that both the efficient markets hypothesis holds, and that the market acts as a risk-neutral representative trader, the price of the contract will be the best estimate of various parameters tied to the probability of that outcome. While these assumptions are clearly too stark, in many settings the biases from violations are quite small, and in practice, the predictions extracted under these assumptions have been shown to be quite reliable.²

Table 1 summarizes three standard types of contracts. These contracts can be used to elicit the probability of a particular outcome occurring, or, if the outcome takes on numerical values, the expected mean or median.

Following the nomenclature introduced in Wolfers and Zitzewitz (2004), a “winner-takes-all” contract costs some amount $\$p$, and pays off, say, $\$100$, if y , initial unemployment claims (in thousands), are reported to be between $y = 330$ and $y = 340$. Under the assumptions above, the price represents the market’s expectation of the probability that initial unemployment claims will be in this range. A family of such contracts with payoffs tied to all likely values of the event can elicit the entire distribution of the market’s beliefs. For example, a

²The literature has considered two possible divergences. When prediction market security outcomes are correlated with a priced asset market factor, such as the overall level of the stock market, prices can be above objective probabilities if a prediction market security offers hedging value. In most applications, differences between risk-neutral and objective probabilities are smaller (see, for example, Snowberg, Wolfers and Zitzewitz, 2007*a*; Wolfers and Zitzewitz, 2008*b*). A second issue is that when beliefs differ, a market may aggregate beliefs in a different way than is statistically optimal. Wolfers and Zitzewitz (2006) find, however, that under most reasonable assumptions about risk-aversion, prices are close to the wealth-weighted mean belief (see also, Manski, 2006; Gjerstad, 2005; Ottaviani and Sørensen, 2010).

Table 1: Contract Types: Estimating Uncertain Quantities or Probabilities

Contract	Example	Details	Reveals market expectation of...
Winner-takes-all	Outcome y : Level of initial unemployment claims (in thousands).	Contract costs $\$p$. Pays $\$1$ if and only if event y occurs. Bid according to the value of $\$p$.	Probability that outcome y occurs.
Index	Contract pays $\$1$ for every 1,000 initial unemployment claims.	Contract pays $\$ y$.	Mean value of outcome y : $\mathbb{E}[y]$.
Spread	Contract pays even money if initial unemployment claims are greater than y^* .	Contract costs $\$1$. Pays $\$2$ if $y > y^*$. Pays $\$0$ otherwise. Bid according to the value of y^* .	Median value of outcome y .

Notes: Adapted from Table 1 in Wolfers and Zitzewitz (2004).

set of contracts, one of which pays off if initial unemployment claims are between 310 and 320, another that pays off if they are between 320 and 330, and so on, recovers the markets' p.d.f. (integrated over bins of 10,000) of initial unemployment claims.³

However, if a researcher is interested in simply obtaining moments of the market's distribution, an "index" contract may be more useful. Such a contract has an unknown payoff tied to the value of the outcome. Continuing the example from above, if y is initial unemployment claims (in thousands), the price for this contract represents the market's expectation of the outcome: $\mathbb{E}[y]$. Higher order moments can be elicited by using multiple index contracts. For example, by adding a contract that pays off the square of initial unemployment claims y^2 , one can recover the variance of the market's beliefs: $\mathbb{E}[y^2] - \mathbb{E}[y]^2$.

Finally, a "spread" contract allows a researcher to elicit a robust measure of the central

³The techniques contained in Bakshi and Madan (2000) and Carr and Madan (2001) may be especially useful in analyzing such a family of contracts.

tendency of the market's beliefs. These contracts, which are familiar to bettors on basketball and American football, feature a fixed cost and fixed payoff if an indicator is below the spread value, however, the spread's value y^* changes as contracts are bought and sold. That is, if the payoff is twice the cost, say \$2 and \$1, respectively, the value of the spread y^* is allowed to vary until an equal number of contracts are bought and sold. The value of the spread y^* is thus the median of the market's distribution. By varying the cost and payoff of the contract, it is possible to elicit other percentiles. For example, a contract that costs \$4 and pays \$5 when $y > y^*$ elicits the value y^* such that the market believes $\text{Prob}(y < y^*) = 0.8$.

We should note that although all of the contracts discussed above have been utilized to some extent, winner-take-all contracts on a single event are by far the most popular form of prediction markets. This is largely because such contracts are straight-forward for market participants to understand. More research is needed to understand the benefits, and limitations, of other forms of contracts and families of those contracts.

2.1 Formats for Administering Prediction Markets

Most prediction markets, like those available on the industry standard Intrade.com, are implemented like equity markets. That is, buyers and sellers interact through a continuous double auction. However, a number of other designs have been suggested, and, in some cases gained traction. The major variations are play-money markets, pari-mutuel markets, and market scoring rules.

Due to concerns about speculation and manipulation, the U.S. Commodities Futures Trading Commission has tightly regulated prediction markets (Arrow et al., 2008). Thus, many markets, especially internal corporate markets, have been established using play money (Bell, 2009). For example, Lumenlogic, a leading provider of prediction markets to businesses, and the Hollywood Stock Exchange, which seeks to forecast entertainment related outcomes, have chosen to run their exchanges using a virtual currency that can either be traded in for prizes, or amassed for prestige. Existing research suggests that these virtual

currency prediction markets can be as accurate as real-money markets, but more research is needed (Servan-Schreiber et al., 2004).

Economic Derivatives, large scale markets run by Goldman Sachs and Deutsche Bank that were tied directly to macroeconomic outcomes, such as initial unemployment claims and non-farm payrolls, were run as occasional, pari-mutuel markets. This structure, familiar from gambling on horse races, allows traders to buy or sell contracts on specific ranges of the economic indicator—similar to the example index contract above. However, the price is not determined until the auction closes—although estimated prices are displayed based on the price that would prevail if no more orders were posted. Once the market is closed and the outcome is realized, pari-mutuel markets split the total amount wagered among those who wagered on the correct range. If, for example, the actual initial unemployment claims (in thousands) were between $y = 330$ and $y = 340$, then each contract on that range would split the total amount wagered on all ranges.⁴ Like the example of index contracts given above, a pari-mutuel market structure will recover the full p.d.f. of the market’s beliefs. However, the pari-mutuel market does this with a single market, rather than the numerous index contracts that would need to be set up.⁵

Finally, market scoring rules (Hanson, 2003, 2007) have shown great promise in the lab, and have been used by Microsoft and other corporations for their internal markets (Abramowicz, 2007). This format starts with a simple scoring rule—which rewards a single person for the accuracy of his or her prediction—and allows other entrants to essentially purchase the right to the reward when they believe they have a more accurate prediction. This format is used largely to reduce speculation and deal with problems that typically arise from thin markets. However, like standard scoring rules, market scoring rules require a market maker to subsidize the reward for accurate predictions. Still, in some situations,

⁴The Economic Derivatives market used a more sophisticated algorithm to determine prices in such a way to maximize trades. For a more detailed description of the design of these markets see Gürkaynak and Wolfers (2005).

⁵Plott (2000) gives an example using a similar structure inside of a large business to produce a full p.d.f. sales forecast.

this may be a small price to pay for the added benefits.

3 Why Prediction Markets Work

Three inter-related facets lead to prediction markets' ability to produce accurate, reliable forecasts. First, the market mechanism is essentially an algorithm for aggregating information. Second, as superior information will produce monetary rewards, there is a financial incentive for truthful revelation. Third, and finally, the existence of a market provides longer-term incentives for specialization in discovering novel information and trading on it. While these facets are inherent in any market, other forecasting mechanisms, such as polling, or employing professional forecasters, lacks one or more of them. For example, polling lacks incentives for truthful revelation, and professional forecasters may have other motivations than simply forecast accuracy (Ottaviani and Sørensen, 2012).

This section documents several examples of how these facets lead to desirable characteristics in practice, such as the rapid incorporation of information, a lack of arbitrage opportunities, and a resistance to manipulation. We then turn to a brief overview of design flaws that may lead markets to fail, and conclude this section with a discussion of more speculative uses of prediction markets. The next section compares prediction market accuracy to other available methods.

As an illustration of the speed at which prediction markets incorporate new information, consider the killing of Osama bin Laden. At 10:25 p.m. Eastern time on May 1st, 2011, Keith Urbahn, the former chief of staff to Defense Secretary Donald Rumsfeld used the social media service Twitter to announce, "So I'm told by a reputable person they have killed Osama Bin Laden. Hot Damn." The first panel of Figure 1 shows the reaction of a prediction market tracking the probability that Osama bin Laden would be captured or killed by December 31st, 2011. Within moments of Urbahn's statement, the price of this contract, which generally saw very little trading, started to rise. Within 25 minutes of this


Figure 1: Information is incorporated quickly, and continuously across time.



So I'm told by a reputable person they have killed Osama Bin Laden. Hot damn.


1 May via [Twitter for BlackBerry](#)

(Donald Rumsfeld's former Chief of Staff)



Notes: Second panel based on data from the IEM, 1988-2000, available at <http://www.biz.uiowa.edu/iem>. Adapted from Wolfers and Zitzewitz (2004).

Figure 2: Prediction markets show little evidence of arbitrage opportunities.



Notes: Data underlying lower line for each data source are bids, data underlying upper line for each data source are asks. Prices collected electronically every four hours by David Pennock, adapted from Wolfers and Zitzewitz (2004).

initial announcement, the probability implied by this contract had risen from 7% to nearly 99%. This final estimate was eight minutes before any mainstream media outlet “broke the story” of bin Laden’s death.

Moreover, prediction markets continuously incorporate new information. The second panel of Figure 1 shows the accuracy of election forecasts from the Iowa Electronic Markets (IEM), an academic-run prediction market based at the University of Iowa, across time. The figure shows quite clearly that as the election approaches, and more information is revealed about likely outcomes, the forecast error decreases steadily.⁶

Prediction markets also evince little evidence of arbitrage opportunities. Contracts within a single exchange are generally arbitrage linked. For example, a rapid increase in the price of a contract tied to the victory of a particular political candidate is generally accompanied

⁶Additionally, the error is generally smaller than polls across time, see Figures 6 and 7, and the surrounding discussion.

by decreases in the prices of contracts linked to other candidates, so the implied probability of *someone* winning an election stays quite close to 100%. Although the general winnowing of prediction market companies in the 2000s makes it more difficult to compare prices across exchanges, Figure 2 shows that in 2003 the prices of different contracts tied to Arnold Schwarzenegger becoming the governor of California exhibited significant variation, but moved in lockstep.⁷

Additionally, prediction markets are quite difficult to manipulate. Attempts by party bosses to manipulate the price of their candidates in turn of the century gambling markets were largely unsuccessful, as were more recent attempts by candidates themselves (Rhode and Strumpf, 2008; Wolfers and Leigh, 2002). More recent, and systematic, attempts at manipulation yielded only brief transitory effects on prices. Rhode and Strumpf (2008) placed random \$500 bets (the largest allowed) on the IEM and found that prices quickly returned to pre-manipulation levels. Camerer (1998) placed large bets in pari-mutuel horse racing markets (which he cancelled moments before the race) to see if this would create a bandwagon effect of follow-on bets. It did not. Finally, evidence from experimental prediction markets, run in the lab, show similar results. In a first experiment where some participants were incentivized to try to manipulate prices, there was little evidence that these participants were successful (Hanson, Oprea and Porter, 2006). A second experiment based incentives on whether or not observers of the market price could be manipulated. Similarly, there was little evidence that manipulators affected the beliefs of observers (Hanson et al., 2011). Indeed, manipulators may increase the accuracy of prediction markets by providing more liquidity (Hanson and Oprea, 2009).

Anecdotal reports of attempts at manipulation abound, and reports indicate they have generally been unsuccessful. One exception we are aware of is manipulation of a contract tied

⁷There was an arbitrage opportunity in 7 of the 105 time periods for which we have data, and the average size the arbitrage opportunity was 1.17 points. However, as positions cannot be moved across exchanges, executing on that opportunity would require holding a long position on one exchange, and a short position on the other until expiry. That is, one must hold \$200 in positions for the duration of the contract in order gain \$1 in arbitrage.

to Hillary Clinton's probability of winning the presidency in 2008, conditional on winning the Democratic nomination. A single, large trader bought contracts on Intrade.com at prices that implied Clinton was much more electable by the general electorate than her Democratic competitors. While the manipulator was able to keep prices high for a significant period of time, the profit opportunity was noticed and discussed by other traders, resulting in huge losses for the large trader. Moreover, the manipulated price garnered only one mention in the mainstream media.⁸

Finally, in most cases, prediction markets seem to satisfy at least the weak form of the efficient markets hypothesis. That is, there is no evidence that trading on past prices can result in a profit. This has been explicitly demonstrated for prediction markets by Leigh, Wolfers and Zitzewitz (2003), Tetlock (2004), and Berg, Forrest and Rietz (2006). In particular, Leigh, Wolfers and Zitzewitz (2003) test prediction markets related to the demise of Sadaam Hussein, and find that an augmented Dickey-Fuller test cannot reject the null that those markets follow a random walk. They also find that a KPSS test rejects the null that prices are trend-stationary. Berg, Forrest and Rietz (2006) also finds that prices in the Iowa Electronic Markets follow a random walk. Finally, Tetlock (2004) finds some evidence of mispricing in prediction markets concerning sporting events, including evidence of over-reaction to news. However, the mispricing is not large enough to allow for profitable trading strategies. Moreover, Tetlock (2004) finds no evidence of mispricing in prediction markets about financial events on the same exchange. More evidence about the efficiency of prediction markets can be found in the large literature on gambling markets. This literature generally finds that betting markets are weakly efficient (discussed in several chapters of Vaughn Williams, 2005; Hausch and Ziemba, 2007).

3.1 Why They (Sometimes) Fail

⁸A summary of this (successful) attempt at manipulation attempt can be found in Zitzewitz (May 30, 2007).


Although prediction markets generally function quite well, design flaws sometimes prevent reliable forecasts. These flaws generally lead to a lack of noise traders (or thin markets) that reduces incentives for discovering, and trading on the basis of, private information (Snowberg, Wolfers and Zitzewitz, 2005). In order to attract noise traders, the subject of a prediction market must be interesting and information must be widely dispersed.

Prediction market contracts must be well specified, so that it is clear when they will (and will not) pay off. However, this specificity may be in tension with making a contract interesting for traders. For example, in 2003, InTrade ran markets that asked, “Will there be a UN Resolution on Iraq (beyond #1441)?” and “Will Saddam be out of office by June 30?” The former is clearly better specified, but the latter had much higher trading volume. Moreover, even the former has some ambiguity: what does it mean for a UN Resolution to be “on” Iraq?

Noise traders may quite rationally choose not to trade in markets where there is a high degree of insider information. For example, despite the high intrinsic interest in who a Supreme Court nominee will be, markets on this topic have routinely failed. This may be due to the fact that most traders are aware that there are very few people with actual information on who the President’s choice will be. This anecdote underlines the importance of prohibiting insider trading: for instance, a market to predict the Institute for Supply Management’s (ISM’s) business confidence measure would be unlikely to function if it were well known that ISM employees were trading in it.

An extreme form of information not being widely dispersed is when there is no information at all to aggregate. For example, prediction markets on whether weapons of mass destruction (WMDs) would be found in Iraq predicted they would very likely be found. The false confidence that could be inspired by such an estimate ignores the fact that there was no information being aggregated by these markets. That is to say, it was unlikely that anyone in Iraq, who might actually have some information (perhaps based on rumors, past experience, or informal discussions with friends and relatives in the government) about whether Iraq’s

Figure 3: The favorite-longshot bias is the most prominent pricing anomaly in sports betting.



Notes: Adapted from Snowberg and Wolfers (2010).

WMD program was likely to exist or not, was trading in these markets.

Finally, it is unclear to what extent prices in prediction markets may be affected by behavioral biases. The most prominent, and well-understood, pricing anomaly in sports betting is the favorite-longshot bias. This pattern, shown across many studies and countries by Figure 3, comes from long-shots being overbet and favorites being underbet relative to their risk-neutral probabilities of victory.⁹ Thus, interpreting the prices of gambles on horses as probabilities will tend to underestimate a horse's probability of victory when that probability is very high, and tend to overestimate a horse's probability of victory when that probability is very low.

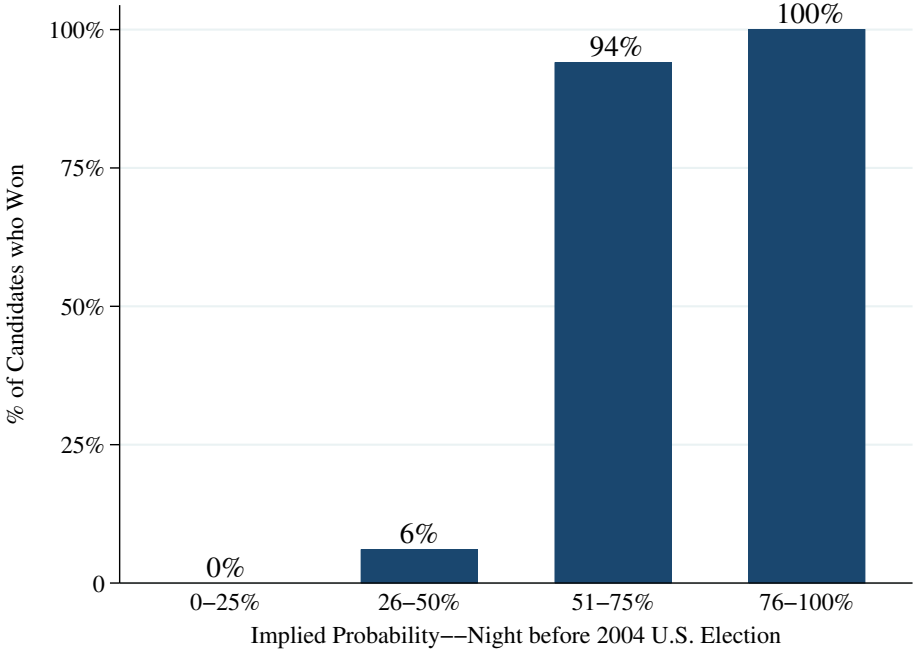
A similar pattern is documented in 2004 U.S. election markets from InTrade.com in Figure 4. Almost all of the contracts that traded above a price of 50 won, and almost all of

⁹As the favorite-longshot bias is a divergence between risk-neutral implied probabilities and actual probabilities, it is also related to pricing phenomena in options markets, see Rubinstein (1994); Tompkins, Ziemba and Hodges (2008).

those that traded below a price of 50 lost, implying that those markets which predicted a high probability of one candidate winning underpredicted the actual probability of victory, while those that predicted a low probability of victory overpredicted the actual probability.

Recent work has attributed the favorite-longshot bias in gambling markets to behavioral phenomena, with Jullien and Salanié (2000) attributing it to asymmetries in the way traders value gains and losses, and Snowberg and Wolfers (2010) attributing it to misperceptions of probabilities. Regardless of the underlying explanation, a good rule of thumb is to use extreme caution when interpreting results based on contracts that imply a risk-neutral probability between 0 and 10%, or 90% and 100%.

Figure 4: The favorite-longshot bias is also apparent in prediction markets.



Notes: Data pools 50 markets on the state-by-state outcome of the electoral college, and 34 markets on U.S. Senate races on Intrade.com (then: Tradesports.com).

3.2 Bleeding Edge Design

The stories of failure above lead to some straight-forward rules for designing prediction markets: make sure the question is well-defined, that there is dispersed information about

the question, and that there is sufficient interest in the question to ensure liquidity. However, such simple guidelines leave much to be desired. For example, what if the question is inherently difficult to define, or there simply is not sufficient interest in the question? These problems have been confronted by corporations using prediction markets, and although there is little academic work on their experiences, we do our best to summarize those experiences here.¹⁰

Corporations are attracted to prediction markets as they can potentially pass unbiased information from a company's front-line employees to senior management.¹¹ However, many questions of interest to executives are not widely interesting, nor are there many employees that have relevant information. This creates a lack of liquidity in markets, perhaps leading to no trading, or, worse, inaccurate predictions. Microsoft has responded to this problem by using a market-making algorithm, which is a variation of the market scoring rules described above (Berg and Proebsting, 2009). The market maker's function in a prediction market is similar to that in any other equity market: it buys and sells when other market participants are not available. However, unlike most market-makers, market-making algorithms are less concerned with making money, and often prefer to lose money if it is likely to increase informational efficiency.

Separately, Hewlett-Packard has used a structure similar to a pari-mutuel market in which a small number of participants are forced to take bets. This implementation is particularly interesting as the market designers use other information about the market participants, such as their risk-attitudes and social connections, to enhance the market's efficiency. For example, if three participants who are known to be friends all bet heavily on the same potential outcome, the market interprets these bets as being based on redundant information,

¹⁰Firms whose internal prediction markets have been mentioned in the public domain include Abbott Labs, Arcelor Mittal, Best Buy, Chrysler, Corning, Electronic Arts, Eli Lilly, Frito Lay, General Electric, Hewlett-Packard, Intel, InterContinental Hotels, Masterfoods, Microsoft, Motorola, Nokia, Pfizer, Qualcomm, Siemens, and TNT (Cowgill, Wolfers and Zitzewitz, 2009). The results of these markets are rarely reported, as they often entail sensitive information.

¹¹Side benefits, such as creating a sense of empowerment and fun among a company's front line employees have also been reported, see Strumpf (2009).

and lowers the impact of each bet on the overall price (Chen and Krakovsky, 2010).

One of the most prominent corporate applications of prediction markets has been to manage R&D portfolios. This is done through “idea markets” which allow employees to buy stock in particular ideas in order to determine which idea(s) are likely to be most profitable for the company. However, the price of the contract depends only on supply and demand, and is not tied to any actual outcome (such as profit if the product is eventually launched). These markets are sometimes described as “preference markets”, and they are similar to pure beauty contests (Marinovic, Ottaviani and Sørensen, 2011). Thus, there is the risk of divergence between the question of interest, and what the market can answer. Indeed, GE, one of the most prominent users of idea markets has found that the originators of ideas trade aggressively to raise the price of their idea and lower the price of other’s ideas (Spears et al., 2009).

Despite the theoretical problems associated with such markets, and the lack of a clear way to evaluate their performance, they are increasingly popular. This has driven a range of experiments with different designs. These experiments vary many of the fine details of market structure, and hence are not easily summarized. Details can be found in Lavoie (2009), Spears et al. (2009), Dahan, Soukhoroukova and Spann (2010). An unusually detailed description of the design and evaluation of a prediction market to help with technology assessment can be found in (Gaspoz, 2011).

A final design constraint is the uncertain legal and regulatory environment surrounding prediction markets (Arrow et al., 2008). The current legal situation is summarized in Bell (2009). Companies are likely best served by consulting with a professional purveyor of prediction market services, a list of which can be found in Berg and Proebsting (2009).

4 Forecast Accuracy

The most notable feature of prediction markets is the accuracy of the forecasts they produce. We illustrate this accuracy through a number of case studies, beginning with an examination of the sole, large-scale, prediction market run on macro-economic outcomes, the Economic Derivatives markets mentioned earlier. We supplement this with case studies from businesses and politics, where prediction markets have been more extensively studied, and arguably, have had a much larger impact.

4.1 Macro Derivatives

For a few years, beginning in October 2002, Goldman Sachs and Deutsche Bank operated markets tied directly to macro-economic outcomes.¹² These markets, described as “Economic Derivatives”, allowed investors to purchase options with payoffs that depended on: growth in non-farm payrolls, retail sales, levels of the Institute for Supply Management’s (ISM’s) manufacturing diffusion index (a measure of business confidence) and initial unemployment claims, among others. The payoffs in these markets were tied to particular data releases, and thus allowed investors to hedge against specific, event-based, risk. The performance of these markets was analyzed by Gürkaynak and Wolfers (2005); this subsection summarizes their main results.


The forecasts gleaned from these markets can be compared with a survey based forecast released by Money Market Services, which typically averages predictions across approximately 30 forecasters, as in Figure 5.¹³ Visual inspection shows that the market-based forecast are weakly more accurate than the survey-based forecast, and this is verified by a numerical analysis.

Table 2 examines both the mean absolute error and mean squared error of both forecasts, normalized by the average forecast error from past surveys. The normalization makes the data across all four series sufficiently comparable to pool all data in the fifth column. For

¹²These markets no longer exist, presumably because the trading volume was not sufficiently high as to make them profitable.

¹³Money Market Services was acquired by Informa in 2003, so some of the survey-based data comes from Action Economics and Bloomberg.

Figure 5: Macro derivatives are weakly more accurate than survey forecasts.



Notes: Adapted from Gürkaynak and Wolfers (2005).

all four series, the market-based forecast is more accurate than the survey-based forecast. Across all series, relying on the market-based forecast would have reduced forecast errors by approximately 5.5% of the average forecast error over the preceding decade. With the small number of observations in each series, the differences are rarely statistically significant, however, the differences in the pooled data are statistically significant at the 5% level.¹⁴

Table 3 examines the forecasting power of each predictor. The first panel reports the correlation of each forecast with actual outcomes. As all of the correlations are quite high, it is clear that both sources have substantial unconditional forecasting power. The middle panel implements Harvey, Leybourne and Newbold (1998) tests of forecast encompassing, and finds that we can reject the null that the market encompasses the survey for Business confidence and initial unemployment claims, we can also reject the null that the survey encompasses the market for business confidence. Additionally, the p-values on the test of

¹⁴Bootstrapped standard errors produce very similar results.

Table 2: Economic derivatives have slightly smaller errors than forecasters.

	Non-farm payrolls	Business confidence (ISM)	Retail sales (ex Autos)	Initial Unemp. Claims	Pooled data
Panel A: Mean Absolute Error					
Economic derivatives	0.723 (.097)	0.498 (.090)	0.919 (.123)	0.645 (.061)	0.680 (.044)
Survey of forecasters	0.743 (.098)	0.585 (.093)	0.972 (.151)	0.665 (.063)	0.719 (.046)
Difference	-0.021 (.034)	-0.087* (.043)	-0.053 (.060)	-0.020 (.023)	-0.039** (.018)
Panel B: Mean Squared Error					
Economic derivatives	0.823 (.240)	0.481 (.257)	1.222 (.262)	0.654 (.126)	0.753 (.088)
Survey of forecasters	0.862 (.268)	0.593 (.296)	1.512 (.364)	0.690 (.130)	0.848 (.114)
Difference	-0.039 (.061)	-0.112 (.082)	-0.289 (.208)	-0.037 (.040)	-0.095** (.018)
Sample Size (Oct. 2002–Jul. 2005)	33	30	26	64	153

Notes: ***, **, * denote statistically significant Diebold and Mariano (1995) / West (1996) tests, implemented as recommended by West (2006), at the 1%, 5% and 10% level. Standard errors in parenthesis. Forecast errors normalized by historical standard error of survey-based forecasts. Adapted from Table 1 in Gürkaynak and Wolfers (2005).

the null that the survey encompasses the market for retail sales and initial unemployment claims are 0.11 and 0.16 respectively. Thus, it is clear that both the market and the survey provide some unique information.

The final panel follows Fair and Shiller (1990) in using a regression-based test of the information content of each forecast. The results are striking: for all four series the coefficient on the market based forecast is statistically indistinguishable from one, and the coefficient

on the survey-based forecast is only statistically different from zero in one series, where it is (perversely) negative. That is, conditioning on the market-based forecast renders the survey-based forecast uninformative. Pooling the data across all four series only reinforces this general pattern.

While the market based forecasts perform better than the survey of forecasters according to most of the statistical criteria examined here, these criteria may not be the most relevant. It is still an open question whether a particular trading strategy, similar to Leitch and Tanner (1991), would fare better or worse using information from forecasters or prediction markets.

4.2 Politics

Prediction markets first gained notoriety for their ability to predict election outcomes. Researchers at the University of Iowa began markets tracking various candidates' vote shares, and chances of victory, in 1988. This experimental market, called the Iowa Electronic Market (IEM), proved to be incredibly accurate. Figure 6 summarizes the predictions of the IEM and compares them with the predictions of polls from Gallup. The figure shows that markets were slightly better predictors than polls the day before the election. However, markets are statistically more accurate than polls in the run-up to the election itself. Over the run-up, markets have half the forecast error of polls (Berg, Forrest and Rietz, 2006; Berg and Rietz, 2006; Berg et al., 2008).

However, as Erikson and Wlezien (2008) notes, there are well-known biases in polls that can and should be controlled for. For example, both parties' candidates usually see their poll numbers climb immediately after their party's convention. Rothschild (2009) compares the forecasts from polls, de-biased in real time, with prediction market prices adjusted for the favorite-longshot bias (discussed above) over the 2008 Presidential election.¹⁵ The results are summarized in Figure 7, which considers prediction markets and polls tied to the outcome

¹⁵To de-bias polls, Rothschild (2009) follows Erikson and Wlezien (2008), which finds the optimal projection from polls at the same point in previous election cycles and eventual elections, and applies that projection to polls from the current electoral cycle.

Table 3: Economic derivatives have slightly smaller errors than forecasters.


	Non-farm payrolls	Business confidence (ISM)	Retail sales (ex Autos)	Initial Unemp. Claims	Pooled data
Panel A: Correlation of Forecast with Actual Outcomes					
Economic derivatives	0.700 (.126)	0.968 (.047)	0.653 (.151)	0.433 (.114)	0.631 (.063)
Survey of forecasters	0.677 (.130)	0.961 (.052)	0.544 (.168)	0.361 (.117)	0.576 (.066)
Panel B: Encompassing Tests (based on absolute error)					
Economic derivatives encompasses survey	-0.048 (.157)	0.194* (.113)	0.123 (.267)	-0.260*** (.089)	-0.062 (.075)
Survey of forecasters encompasses market	0.081 (.159)	-0.269** (.110)	-0.471 (.326)	-0.125 (.089)	0.096 (.081)
Panel C: Horse Race Regression (Fair-Shiller)					
$Actual_t = \alpha + \beta * Economic\ Derivatives_t + \gamma * Survey\ Forecast_t (+survey\ fixed\ effects)$					
Economic derivatives	1.06 (.78)	0.91** (.37)	1.99** (.79)	1.64*** (.60)	1.25*** (.29)
Survey of forecasters	-0.14 (.89)	0.17 (.38)	-1.03 (1.10)	-1.21* (.68)	-0.24 (.30)
Adjusted R^2	0.46	0.93	0.40	0.20	0.99
Sample Size (Oct. 2002–Jul. 2005)	33	30	26	64	153

Notes: ***, **, * denote statistically significant coefficients at the 1%, 5% and 10% level with standard errors in parenthesis. Forecast errors normalized by historical standard error of survey-based forecasts. Adapted from Table 1 in Gürkaynak and Wolfers (2005).

of the presidential election in each of the 50 states.

As the figure shows, the mean squared error of polls is much higher, sometimes statistically significantly so, over the first 70 days of the election cycle. For the second half, polls and markets switch places (twice), and one is not clearly better than the other. However, as shown in Table 4, in Fair and Shiller (1990) type horse-race regressions, prediction mar-

Figure 6: Prediction markets are more accurate even the night before an election.



Notes: Market forecast is closing price on election eve; Gallup forecast is final pre-election projection.

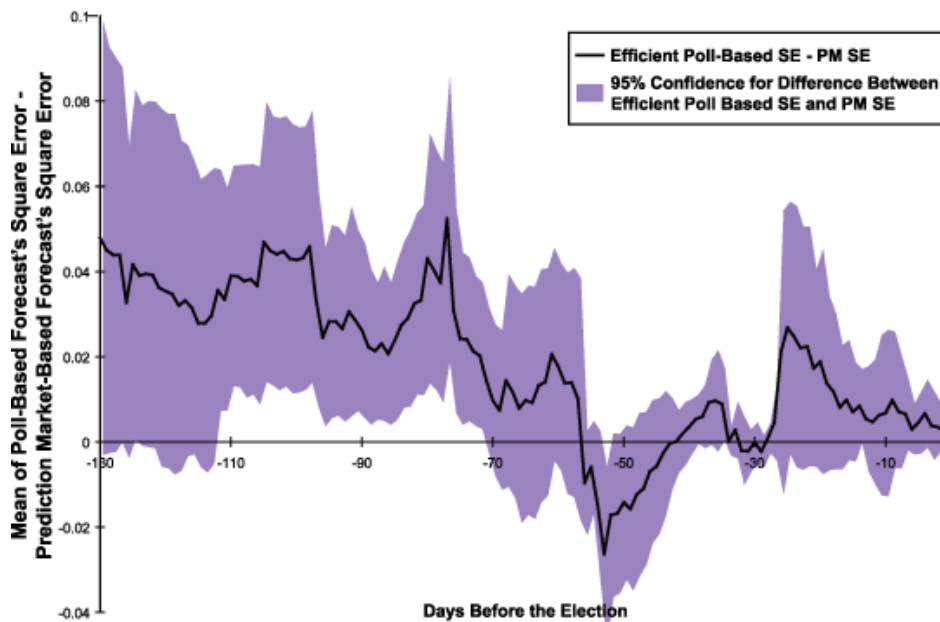
markets encapsulate all of the information in polls when contracts that indicate a probability of 90% of one or the other candidate winning that state are dropped from the sample. When these markets are included, the coefficient on the forecast from the prediction market is still statistically significant and close to one, but the coefficient on de-biased polls is statistically different from zero. Taking all of this evidence together, raw prediction market prices provide superior forecasts to raw polls, but polls may contain some additional information, especially in races where one candidate dominates.

4.3 Business

Corporations have aggressively used prediction markets to help with their internal forecasts. As noted in Section 3.2, there is little academic evidence of their impact. However, three important exceptions deserve a mention here.

First, in a seminal study, Chen and Plott (2002) ran eight prediction markets within

Figure 7: Prediction markets are generally more accurate than polls, even after removing known biases.



Notes: Each datapoint is the difference in mean-squared error between two different types of forecast over the 50 electoral college races in 2008. Each forecast is generated by taking raw data from a poll or prediction market and utilizing the most efficient transformation from raw data into forecasts, as outlined in Rothschild (2009).

Table 4: Comparing the information in prediction markets and polls.

Horse Race Regression (Fair-Shiller)			
$Vote\ Share = \alpha + \beta * Poll + \gamma * Market\ Prices$			
	All Observations	Before Labor Day	Less-certain Races
De-biased Gallup Poll	0.296*** (.060)	0.458*** (.100)	0.195 (.164)
De-biased Prediction Market Prices	0.759*** (.072)	0.754*** (.099)	0.788*** (.099)
N	8,361	4,354	3,167

Notes: ***, **, * denote statistical significance at the 1%, 5% and 10% level with robust standard errors in parenthesis. Table uses the data from Rothschild (2009), with prediction market prices converted into expected vote shares using past trends. Less-certain races are those where neither candidate has a (projected) 90% chance of winning. More details of the procedure for de-biasing polls and prediction market prices can be found in Rothschild (2009).

Hewlett-Packard to forecast important variables like quarterly printer sales. These results showed that the markets were more accurate than the company's official forecasts. This improvement in accuracy was obtained even though the markets had closed, and final prices were known, at the time the official forecast was made.

Second, Cowgill, Wolfers and Zitzewitz (2009) analyze data from 270 markets run inside of Google. While the markets were, in general, quite accurate, and often provided forecasts that would have been difficult to obtain in anything other than an ad hoc way, there were identifiable biases in some market prices. In particular, when markets involved the performance of Google as a company, optimistic outcomes were forecast to be more likely to happen than they actually were. Moreover, this "optimistic bias" was more pronounced among traders who were newer employees, and on days when Google's stock appreciated.

Third, and finally, Berg, Neumann and Rietz (2009) ran several prediction markets to predict Google's market cap at the end of the first day of trading. Notably, these prediction could be compared to the auction that Google used in setting its IPO price. The prediction market fared quite well: its prediction was 4% above the actual market cap, while the IPO price was 15% below. Had the company set its IPO price based on the prediction market price, they would have earned \$225 million more in their IPO.

5 Discovering Economic Models

The information from prediction markets has also proven quite useful in augmenting event studies (Snowberg, Wolfers and Zitzewitz, 2011). This use of prediction markets may be of particular interest to economic forecasters, as they reveal pieces of the economic model underlying the markets' reaction to various types of information. We show first how prediction markets can be used to measure the uncertainty surrounding an event, before exploring two examples of prediction market event studies. The first example measures the expected impact of the second Iraq war on oil prices, while the second example examines the impact

of politics on broad stock market indices.


5.1 Disagreement and Uncertainty

Consumers of economic forecasts care not just about the mean prediction, but the uncertainty attached to it as well. As professional forecasters generally only give mean forecasts, it has become common practice to report *disagreement*—the standard error of economic forecasters’ forecasts—as a proxy for uncertainty. This practice is theoretically founded: Laster, Bennett and Geoum (1999) shows that, under reasonable assumptions, the distribution of forecasts will match the distribution of beliefs. However, data from prediction markets shows that disagreement between forecasters is a poor proxy for uncertainty.

Markets, such as the economic derivative markets above, allow one to recover the entire distribution of the markets beliefs. In turn, this may be used to calculate the *uncertainty*—the standard deviation—of market based forecasts. Gürkaynak and Wolfers (2005) shows that this measure is very close to the RMSE of the market based (mean) forecast, so the market’s uncertainty seems to be well-calibrated. However, Figure 8 shows that disagreement tends to be much smaller than actual uncertainty.

Moreover, there is generally a poor correlation between disagreement and uncertainty. Table 5 regresses uncertainty on disagreement. Panel A shows a statistically significant, positive correlation between the two measures for all series except ISM. The fifth column shows that the coefficients are jointly quite significant, suggesting a strong, contemporaneous relationship. Panel B focuses on lower-frequency variation. There is still a correlation between the two series, but it is substantially lower. Note that the correlation is likely overstated as standard errors here are not corrected for the autocorrelation generated by smoothing. The joint test now fails to achieve statistical significance at conventional levels. Together, these results suggest that disagreement is poorly calibrated to actual forecast error, and that it is a poor proxy for uncertainty.

Figure 8: Uncertainty is generally greater than disagreement.



Notes: Dashed lines show 5-period, centered, moving averages. Adapted from Gürkaynak and Wolfers (2005).

5.2 War

Wars can have large economic impacts, and, in the post WWII era, usually have a fairly long buildup. However, as each war is unique, it is difficult to build wars into economic forecasts, even when the probability of war is very high. In particular, it is quite difficult to make conditional forecasts in anything other than an ad hoc way, as much of the information available may be tainted by political considerations, and, moreover, such a large-scale event may have general equilibrium effects that are poorly understood.

Prediction market event studies present the possibility of improving economic forecasts when such a unique event becomes likely. By tracking the probability of such an event through a prediction market, and correlating movements in the prediction market with asset prices, it becomes possible to understand the market's assessment of how such an event will affect outcomes such as the price of oil, or the returns to securities.

Leigh, Wolfers and Zitzewitz (2003) performed a prediction market event study in the

Table 5: Disagreement among forecasters is a poor proxy for uncertainty.

	Non-farm payrolls	Business confidence (ISM)	Retail sales (ex Autos)	Joint Significance (F-test)	
Panel A: Contemporaneous Relationship					
$Uncertainty_t = \alpha + \beta * Disagreement_t$					
Disagreement	0.66** (.29)	-0.03 (.12)	0.44** (.16)	0.27*** (.07)	
Constant	73.6*** (10.39)	2.04*** (.13)	0.36*** (.03)	10.86*** (.47)	$p = .0002$
Adjusted R^2	0.11	-0.03	0.20	0.17	
Panel B: Low Frequency—5 Period, Centered, Moving Averages					
$Smoothed\ Uncertainty_t = \alpha + \beta * Smoothed\ Disagreement_t$					
Disagreement	0.55 (.47)	0.10 (.10)	0.65** (.24)	.32** (.06)	
Constant	77.7*** (16.8)	1.89*** (.11)	0.32*** (.05)	10.5*** (.37)	$p = .15$
Adjusted R^2	0.01	-0.002	0.23	0.32	

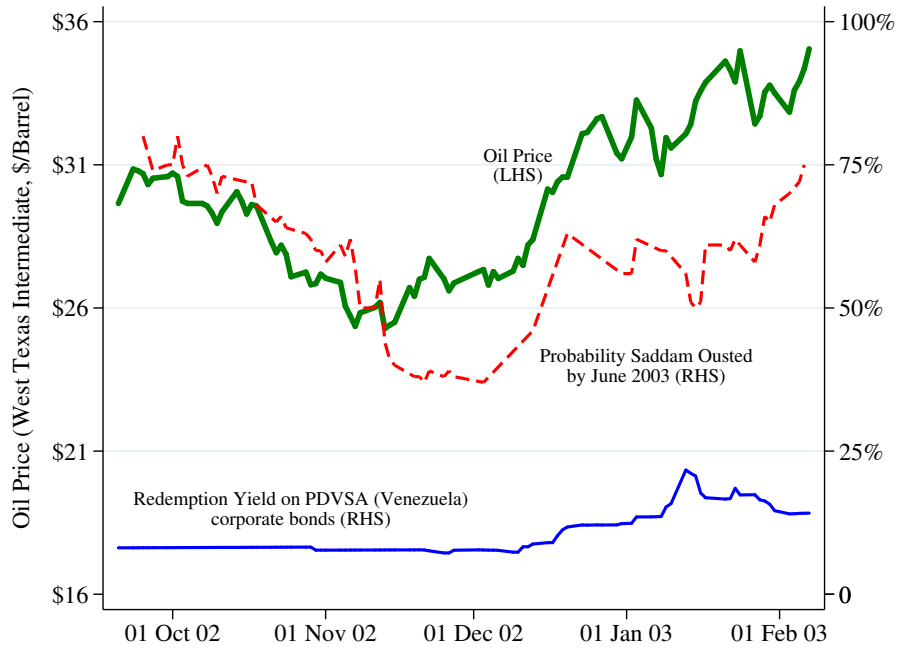
Notes: ***, **, * denote statistically significant regression coefficients at the 1%, 5% and 10% level with standard errors in parenthesis. Adapted from Table 5 in Gürkaynak and Wolfers (2005).

buildup to the second Iraq war. Its prospective analysis showed that every 10% increase in the probability of war lead to a \$1 increase in the spot price of oil. As shown in Figure 9, they used a prediction market, run by Intrade.com, that measured the probability that Saddam Hussein would be out of office by a certain date as the probability of war a few months before that date.¹⁶ As can be seen from the figure, as the probability of war fluctuated, so did the spot price of oil.

Moreover, the S&P 500 dropped 1.5% for every 10% increase in the probability of war. This implied that the market anticipated that a war would decrease stock prices by 15


¹⁶This is a reasonable proxy—out of office was defined as no longer controlling the center of Baghdad.

Figure 9: Response of oil prices to the probability of war.



Notes: Adapted from Leigh, Wolfers and Zitzewitz (2003).

Figure 10: Predicted distributions of the S&P 500 for various probabilities of war.



Notes: Adapted from Wolfers and Zitzewitz (2008b).

percentage points. A further analysis of S&P 500 options revealed that the distribution of potential outcomes was quite negatively skewed. Option prices suggested there was a 70% probability that war would have a moderately negative effect of 0 to -15 percentage points, a 20% probability of a -15 to -30 percentage-point effect, and a 10% chance of even larger declines. Figure 10 presents this information in a slightly disaggregated way: for each of four probabilities of war, it shows the full state-price distribution. It is clear that in the case of certain war (the thick, solid line), the mode of the state-price distribution is substantially shifted from what it would be under a zero chance of war (the thin, dashed line). Moreover, in certain war, the markets predicted a non-negligible probability that the S&P 500 would fall below 500 (Wolfers and Zitzewitz, 2008*b*).

The striking magnitude of this effect stands in contrast with the conventional wisdom, based in part on studies such as Cutler, Poterba and Summers (1989), that political and military news explain only a small portion of market movements. In contrast, Wolfers and Zitzewitz (2008*b*) finds that, “[O]ver 30% of the variation in the S&P and 75% of the variation in spot oil prices between September 2002 and February 2003 can be explained econometrically by changes in the probability of war (and, for oil, the Venezuelan crisis).” The authors note that it is possible to reconcile this result with the conventional wisdom by noting that typically when economically important events happen, they are often near certainties. That is to say, if the actual declaration of war comes when the market’s assessment of the probability of war is already 95%, then any correlated market movements must be scaled by a factor of 20 to appreciate the full economic magnitude of the event. However, determining the markets’ assessment of the probability of war immediately preceding an event is quite difficult without prediction markets.

Additionally, prediction markets can be constructed to measure the overall economic cost of a military intervention. For example, to determine the state price distribution of the S&P 500, a researcher could combine S&P 500 futures with a prediction market on the probability of war, as shown in Figure 10. A cleaner design would be to issue options on the S&P 500

that pay off only if the U.S. has gone to war, with all transactions reversed if the U.S. does not go to war. By combining these options with prediction markets tied directly to whether or not the U.S. goes to war, and, for example, the level of troop deployments in the case of war, one could recover state price distributions under different military scenarios. This (largely theoretical) use of prediction markets has been described as “decision markets” or, more spectacularly, “futarchy” (Hanson, 1999; Berg and Rietz, 2003; Hahn and Tetlock, 2005; Hanson, 2011).¹⁷

5.3 Politics

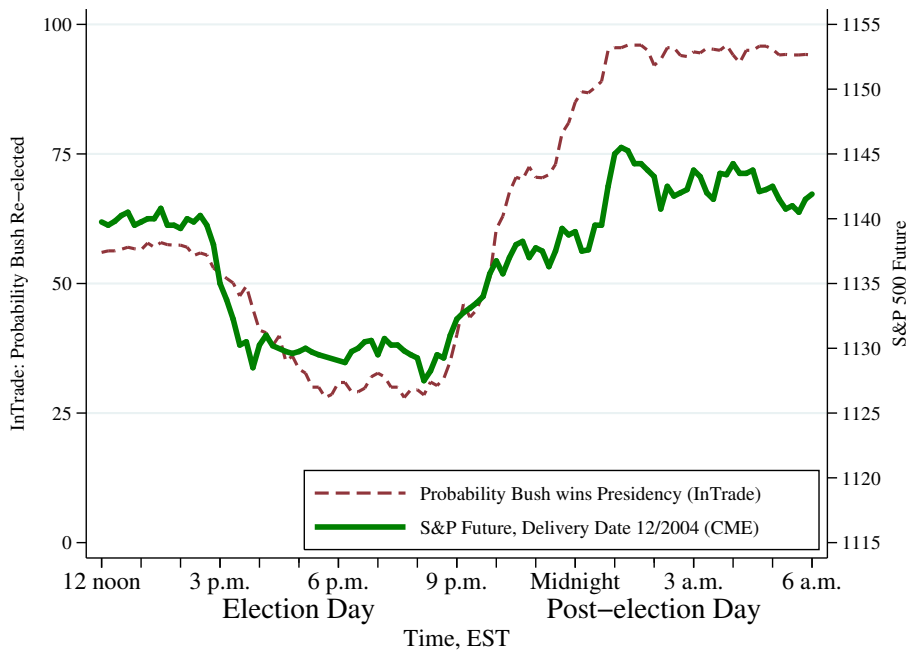
Prediction markets can also help to determine the effect of more routine political events, such as elections. A large literature in political science, economics, and the popular press argues about the effect of various candidates and parties, and the policies they endorse, on the economy. Despite some evidence that broad stock market indices in the U.S. may perform up to 20% better under Democrats than Republicans (Santa-Clara and Valkanov, 2003), political considerations are rarely reflected in economic forecasts.

This is at least partly due to the fact that there is no academic consensus on even the direction of the effect of political outcomes on economic variables. However, markets exhibit much more consistency. In particular, by using a prediction market event study on election night 2004, Snowberg, Wolfers and Zitzewitz (2007*a*) shows that broad stock market indices rose approximately 2% on news of Bush’s victory over Kerry. Moreover, this effect seems to be relatively consistent across time, as, using data from 1880-2004, the same authors show that a Republican victory caused a broad-based market index to rise by approximately 2.5%.

Figure 11 shows the price of a prediction market that paid off if Bush won in 2004, expressed as risk-neutral probabilities, and the value of a near-month S&P 500 future over election night 2004. The prices of these two securities track each other quite closely. The probability of Bush’s re-election starts near 55%, and declines by about 30% on flawed exit

¹⁷Malhotra and Snowberg (2009) provide an example of how to use a decision market to pick the presidential candidate that would maximize a particular party’s chance of winning.

Figure 11: Prediction markets uncover the market’s reaction to political changes.



Notes: Adapted from Snowberg, Wolfers and Zitzewitz (2007a).

polls showing Kerry with a lead in some swing states. At approximately the same time, the S&P 500 future decreases in value by approximately 0.7%. Thus, this particular event study indicates that a Bush victory would increase the value of the S&P 500 by $0.7\%/30\% = 2.3\%$. As early vote totals were released, showing the faults of the earlier poll results, Bush’s probability of re-election climbed 65%, and the S&P rose by about 1.3%. Thus, this event study indicates that a Bush victory would increase the S&P 500 by $1.3\%/65\% = 2\%$.

A first differences regression essentially averages together a large number of event studies. In particular, estimates of

$$\Delta(\log(\text{Financial Variable}_t)) = \alpha + \beta\Delta(\text{Re - election Probability}_t) + \varepsilon_t$$

are shown for a number of different financial variables in Table 6. Estimates based on 10-minute and 30-minute differences are consistent, although the results based on 30-minute differences have slightly larger coefficients, reflecting the fact that the prediction market was

Table 6: High-frequency data from prediction markets allows for precise estimates of the effect of elections on economic variables.

		10-Minute First Differences		30-Minute First Differences	
		Estimated Effect of Bush Presidency	n	Estimated Effect of Bush Presidency	n
Dependent Variable: $\Delta(\text{Log}(\text{Financial Index}))$					
S&P 500		0.015*** (.004)	104	0.021*** (.005)	35
Dow Jones Industrial Average		0.014*** (.005)	84	0.021*** (.006)	29
Nasdaq 100		0.017*** (.006)	104	0.024*** (.008)	35
U.S. Dollar (vs. Trade-weighted basket)		0.004 (.003)	93	0.005** (.003)	34
Dependent Variable: $\Delta(\text{Price})$					
Light Crude Oil Futures	December '04	1.110*** (.371)	88	1.706*** (.659)	29
	December '05	0.652* (.375)	85	1.020 (.610)	28
	December '06	-0.580 (.783)	63	-0.666 (.863)	21
Dependent Variable: $\Delta(\text{Yield})$					
2-Year T-Note Future		0.104* (.058)	84	0.108*** (.036)	30
10-Year T-Note Future		0.112** (.050)	91	0.120** (.046)	31

Notes: ***, **, * denote statistically significant regression coefficients at the 1%, 5% and 10% level with robust standard errors in parenthesis. The sample period is noon Eastern Time on 11/2/2004 to six a.m. on 11/3/2004. Election probabilities are the most recent transaction prices, collected every ten minutes from InTrade.com (then TradeSports.com), S&P, Nasdaq, and foreign exchange futures are from the Chicago Mercantile Exchange; Dow and bond futures are from the Chicago Board of Trade, while oil futures data are from the New York Mercantile Exchange. Equity, bond and currency futures have December 2004 delivery dates. Yields are calculated for the Treasury futures using the daily yields reported by the Federal Reserve for 2- and 10-year Treasuries, and projecting forward and backward from the bond market close at 3 p.m. using future price changes and the future's durations of 1.96 and 7.97 reported by CBOT. The trade-weighted currency portfolio includes six currencies with CME-traded futures (the Euro, Yen, Pound, Australian and Canadian Dollars, and the Swiss Franc). Adapted from Table 1 in Snowberg, Wolfers and Zitzewitz (2007a).

slower to incorporate information than financial markets, as is apparent from Figure 11. As can be seen from the table, a Bush victory increased all three major US stock indices by 2–2.5%. Consistent with expectations of expansionary fiscal policies, the Dollar and Bond Yields rose, as did near-term expectations of the price of oil.

Were this just a one-off result, it would make little sense to add political changes to a forecasting model. However, as Figure 12 and Table 7 show, Republican elections routinely lead to an increase of about 2.5% in a broad equity market index. The first panel of the figure, and first column of the table, show the relationship between the change in probability of a Republican President over election night and the change in a broad-based market index from market close the day before the election, to market close the day after.¹⁸ That is, the table contains an estimate of:

$$\Delta(\text{Market Index}_t) = \alpha + \beta\Delta(\text{Republican President}_t) + \varepsilon_t$$


where $\Delta(\text{Republican President}_t) = \mathbb{I}(\text{Republican President}_t) - P(\text{Republican Election}_t)$

where $\mathbb{I}(\text{Republican Elected}_t)$ takes a value of one if a Republican was elected in year t , and zero otherwise, and $P(\text{Republican Election}_t)$ is the expected probability of a Republican victory, according to the price of an historical prediction market (from Rhode and Strumpf, 2004) the night before the election.

The second panel of the figure, and second column of the table, show the role of prediction markets in this estimate. Lacking prediction market data, Santa-Clara and Valkanov (2003) simply fix $P(\text{Republican Election}_t) = 0.5$ for all years from which they have data: 1928–1996. This results in the smaller, statistically insignificant estimate found in the second column of Table 7. The third column of the same table contains data from all years available in Snowberg, Wolfers and Zitzewitz (2007*a*), and shows a precisely estimated, stable effect of

¹⁸Equity index values are from Schwert's (1990) daily equity returns data, which attempts to replicate returns on a value-weighted total return index, supplemented by returns on the CRSP-value-weighted portfolio since 1962. The prediction market prices come from the curb exchange on Wall Street, where exchanges on various political candidates ran up until shortly after WWII (Rhode and Strumpf, 2004).

Figure 12: Prediction markets uncover the market's reaction to political changes.



Notes: Adapted from Snowberg, Wolfers and Zitzewitz (2011).

Table 7: Prediction markets identify the effect of elections on markets.

	Dependent Variable: Stock returns from election-eve close to post-election close		
$\Delta P(\text{Republican President})$ (<i>From Prediction Markets</i>)	0.0297** (.118)		0.0249*** (.0082)
$\mathbb{I}(\text{Republican President}) - 0.5$ (<i>As in Santa-Clara and Valkanov</i>)		0.0128 (.0089)	
Constant	-0.0102 (.0059)	-0.0027 (0.0040)	0.0014 (0.0028)
Sample	1928–1996	1928–1996	1880–2004
N	18	18	32

Notes: ***, **, * denote statistical significance at the 1%, 5% and 10% level with robust standard errors in parenthesis. Adapted from Table 5 in Snowberg, Wolfers and Zitzewitz (2007a).

the election of a Republican president of about 2.5%. The contribution of prediction markets is clear: a more precise estimate of the probability of one or the other candidate winning allows for a better estimate of the effect of political candidates on financial markets.

6 Conclusion

Over the last half decade, many economists, and the public, have re-evaluated the efficiency of financial markets. While these re-evaluations have not been favorable to markets, it is important to keep in mind the alternative (Zingales, 2010). In the case of forecasting, the alternative is often professional forecasters, polls, pundits, or a combination of the three. As we have shown in these paper, prediction markets out-perform both professional forecasters and polls in a variety of statistical tests.

We have shown that prediction markets have many of the properties expected under the efficient markets hypothesis. In particular, they are difficult to manipulate, lack significant arbitrage opportunities, aggregate information quickly and in a seemingly efficient manner.

Evidence of efficiency can be seen in the macro-derivatives markets, which out-perform professional forecasters, or in political prediction markets, which out-perform polls.

However, prediction markets are not a panacea. In particular, care must be taken when designing prediction markets to ensure they are interesting, well-specified, and are not subject to excessive insider information. More pernicious problems come from behavioral biases, such as those underlying the favorite-longshot bias, and knowing when there is dispersed information that can be aggregated.

With that said, we believe the real promise of prediction markets comes not from their ability to predict particular events. Rather, the real promise lies in using these markets, often several at a time, to test particular economic models, and use these models to improve economic forecasts.

References

- Abramowicz, Michael. 2007. "The Hidden Beauty of the Quadratic Market Scoring Rule: A Uniform Liquidity Market Maker, with Variations." *The Journal of Prediction Markets* 1(2):111–125.
- Alquist, Ron and Robert J. Vigfusson. 2012. Forecasting Oil Prices. In *Handbook of Economic Forecasting, Volume 2*. Elsevier: Handbooks in Economics series.
- Arrow, Kenneth J., Robert Forsythe, Michael Gorham, Robert Hahn, Robin Hanson, John O. Ledyard, Saul Levmore, Robert Litan, Paul Milgrom, Forrest D. Nelson, George R. Neumann, Marco Ottaviani, Thomas C. Schelling, Robert J. Shiller, Vernon L. Smith, Erik Snowberg, Cass R. Sunstein, Paul C. Tetlock, Philip E. Tetlock, Hal R. Varian, Justin Wolfers and Eric Zitzewitz. 2008. "The Promise of Prediction Markets." *Science* 320(5878):877–878.
- Bakshi, Gurdip and Dilip Madan. 2000. "Spanning and Derivative-Security Valuation." *Journal of Financial Economics* 55(2):205–238.
- Bell, Tom W. 2009. "Private Prediction Markets and the Law." *The Journal of Prediction Markets* 3(1):89–110.
- Berg, Henry and Todd A. Proebsting. 2009. "Hanson's Automated Market Maker." *The Journal of Prediction Markets* 3(1):45–59.
- Berg, Joyce E., George R. Neumann and Thomas A. Rietz. 2009. "Searching for Google's Value: Using Prediction Markets to Forecast Market Capitalization Prior to an Initial Public Offering." *Management Science* 55(3):348–361.
- Berg, Joyce E. and Thomas A. Rietz. 2003. "Prediction Markets as Decision Support Systems." *Information Systems Frontiers* 5(1):79–93.
- Berg, Joyce E. and Thomas A. Rietz. 2006. The Iowa Electronic Market: Lessons Learned and Answers Yearned. In *Information Markets: A New Way of Making Decisions in the Public and Private Sectors*, ed. Robert Hahn and Paul Tetlock. Washington, D.C.: AEI-Brookings Joint Center.
- Berg, Joyce, Nelson Forrest and Thomas Rietz. 2006. "Accuracy and Forecast Standard Error in Prediction Markets." University of Iowa, *mimeo*.
- Berg, Joyce, Robert Forsythe, Forrest Nelson and Thomas Rietz. 2008. Results from a Dozen Years of Election Futures Markets Research. In *The Handbook of Experimental Economics Results*, ed. Charles R. Plott and Vernon L. Smith. Elsevier: Handbooks in Economics series.
- Camerer, Colin F. 1998. "Can Asset Markets be Manipulated? A Field Experiment with Racetrack Betting." *Journal of Political Economy* 106(3):457–482.

- Carr, Peter and Dilip Madan. 2001. “Optimal Positioning in Derivative Securities.” *Quantitative Finance* 1(1):19–37.
- Chen, Kay-Yut and Charles R. Plott. 2002. “Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem.” California Institute of Technology Social Science Working Paper #1131.
- Chen, Kay-Yut and Marina Krakovsky. 2010. *Secrets of the Moneylab: How Behavioral Economics Can Improve Your Business*. Portfolio / Penguin.
- Cowgill, Bo, Justin Wolfers and Eric Zitzewitz. 2009. “Using Prediction Markets to Track Information Flows: Evidence from Google.” Dartmouth College, *mimeo*.
- Cutler, David M., James M. Poterba and Lawrence H. Summers. 1989. “What Moves Stock Prices?” *Journal of Portfolio Management* 15(3):4–12.
- Dahan, Ely, Arina Soukhoroukova and Martin Spann. 2010. “New Product Development 2.0: Preference Markets—How Scalable Securities Markets Identify Winning Product Concepts and Attributes.” *Journal of Product Innovation Management* 27(7):937–954.
- Diebold, Francis X. and Robert S. Mariano. 1995. “Comparing Predictive Accuracy.” *Journal of Business and Economic Statistics* 13(3):134–144.
- Duffee, Gregory. 2012. Forecasting Interest Rates. In *Handbook of Economic Forecasting, Volume 2*, ed. Allan Timmermann and Graham Elliott. Elsevier: Handbooks in Economics series.
- Erikson, Robert S. and Christopher Wlezien. 2008. “Are Political Markets Really Superior to Polls as Election Predictors?” *Public Opinion Quarterly* 72(2):190–215.
- Fair, Ray C. and Robert J. Shiller. 1990. “Comparing Information in Forecasts from Econometric Models.” *American Economic Review* 80(3):375–389.
- Gaspoz, Cédric. 2011. *Prediction Markets Supporting Technology Assessment*. Self-published.
- Gjerstad, Steven. 2005. “Risk Aversion, Beliefs, and Prediction Market Equilibrium.” University of Arizona, *mimeo*.
- Gürkaynak, Refet and Justin Wolfers. 2005. Macroeconomic Derivatives: An Initial Analysis of Market-Based Macro Forecasts, Uncertainty, and Risk. In *NBER International Seminar on Macroeconomics 2005*, ed. Jeffrey A. Frankel and Christopher A. Pissarides. Cambridge, MA: MIT Press pp. 11–50.
- Hahn, Robert W. and Paul C. Tetlock. 2005. “Using Information Markets to Improve Public Decision Making.” *Harvard Journal of Law and Public Policy* 28:213–289.
- Hanson, R. 2007. “Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation.” *The Journal of Prediction Markets* 1(1):3–15.
- Hanson, Robin. 1999. “Decision Markets.” *IEEE Intelligent Systems* 14(3):16–19.

- Hanson, Robin. 2003. "Combinatorial Information Market Design." *Information Systems Frontiers* 5(1):107–119.
- Hanson, Robin. 2011. "Shall we Vote on Values, but Bet on Beliefs?" *Journal of Political Philosophy* forthcoming.
- Hanson, Robin, Ryan Oprea, David Porter, Chris Hibbert and Dorina Tila. 2011. "Can Manipulators Mislead Prediction Market Observers?" George Mason University, *mimeo*.
- Hanson, Robin and Ryan Oprea. 2009. "A Manipulator Can Aid Prediction Market Accuracy." *Economica* 76(302):304–314.
- Hanson, Robin, Ryan Oprea and David Porter. 2006. "Information Aggregation and Manipulation in an Experimental Market." *Journal of Economic Behavior & Organization* 60(4):449–459.
- Harvey, David I., Stephen J. Leybourne and Paul Newbold. 1998. "Tests for Forecast Encompassing." *Journal of Business & Economic Statistics* 16(2):254–259.
- Hausch, Donald and William Ziemba, eds. 2007. *Handbook of Sports and Lottery Markets*. Elsevier: Handbooks in Finance series.
- Jullien, Bruno and Bernard Salanié. 2000. "Estimating Preferences Under Risk: The Case of Racetrack Bettors." *Journal of Political Economy* 108(3):503–530.
- Laster, D., P. Bennett and In Sun Geoum. 1999. "Rational Bias in Macroeconomic Forecasts." *Quarterly Journal of Economics* 114(1):293–318.
- Lavoie, Jim. 2009. "The Innovation Engine at Rite-Solutions: Lessons From the CEO." *Journal of Prediction Markets* 3(1):1–11.
- Leigh, Andrew, Justin Wolfers and Eric Zitzewitz. 2003. "What do Financial Markets Think of the War with Iraq?" *NBER Working Paper #9587*.
- Leitch, Gordon and J. Ernest Tanner. 1991. "Economic Forecast Evaluation: Profits Versus the Conventional Error Measures." *The American Economic Review* 81(3):580–590.
- Malhotra, Neil and Erik Snowberg. 2009. "The 2008 Presidential Primaries through the Lens of Prediction Markets." California Institute of Technology, *mimeo*.
- Manski, Charles F. 2006. "Interpreting the Predictions of Prediction Markets." *Economics Letters* 91(3):425–429.
- Marinovic, Iván, Marco Ottaviani and Peter Norman Sørensen. 2011. Modeling Idea Markets: Between Beauty Contests and Prediction Markets. In *Prediction Markets*, ed. Leighton Vaughn Williams. Routledge.
- Ottaviani, Marco and Peter N. Sørensen. 2010. "Aggregation of Information and Beliefs: Asset Pricing Lessons from Prediction Markets." Northwestern University, *mimeo*.

- Ottaviani, Marco and Peter Norman Sørensen. 2012. Forecasters' Objectives and Strategies. In *Handbook of Economic Forecasting, Volume 2*. Elsevier: Handbooks in Economics series.
- Plott, Charles R. 2000. "Markets as Information Gathering Tools." *Southern Economic Journal* 67(1):2–15.
- Plott, Charles R. and Shyam Sunder. 1982. "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models." *The Journal of Political Economy* 90(4):663–698.
- Plott, Charles R. and Shyam Sunder. 1988. "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets." *Econometrica* pp. 1085–1118.
- Rhode, Paul W. and Koleman S. Strumpf. 2004. "Historical Presidential Betting Markets." *Journal of Economic Perspectives* 18(2):127–142.
- Rhode, Paul W. and Koleman S. Strumpf. 2008. "Historical Election Betting Markets: An International Perspective." University of Arizona, *mimeo*.
- Rothschild, David. 2009. "Forecasting Elections Comparing Prediction Markets, Polls, and Their Biases." *Public Opinion Quarterly* 73(5):895–916.
- Rubinstein, Mark. 1994. "Implied Binomial Trees." *Journal of Finance* 69(3):771–818.
- Santa-Clara, Pedro and Rossen Valkanov. 2003. "The Presidential Puzzle: Political Cycles and the Stock Market." *The Journal of Finance* 58(5):1841–1872.
- Schwert, G. William. 1990. "Indexes of U.S. Stock Prices from 1802-1987." *Journal of Business* 63(3):399–426.
- Servan-Schreiber, Emile, Justin Wolfers, David M. Pennock and Brian Galebach. 2004. "Prediction Markets: Does Money Matter?" *Electronic Markets* 14(3):243–251.
- Snowberg, Erik and Justin Wolfers. 2010. "Explaining the Favorite-Longshot Bias: Is it Risk-Love or Misperceptions?" *Journal of Political Economy* 118(4):723–746.
- Snowberg, Erik, Justin Wolfers and Eric Zitzewitz. 2005. Information (In)efficiency in Prediction Markets. In *Information Efficiency in Betting Markets*, ed. Leighton Vaughn Williams. Cambridge University Press pp. 366–386.
- Snowberg, Erik, Justin Wolfers and Eric Zitzewitz. 2007a. "Partisan Impacts on the Economy: Evidence from Prediction Markets and Close Elections." *The Quarterly Journal of Economics* 122(2):807–829.
- Snowberg, Erik, Justin Wolfers and Eric Zitzewitz. 2007b. Prediction Markets: From Politics to Business (and Back). In *Handbook of Sports and Lottery Markets*, ed. Donald Hausch and William Ziemba. Elsevier: Handbooks in Finance series pp. 385–402.

- Snowberg, Erik, Justin Wolfers and Eric Zitzewitz. 2011. How Prediction Markets can Save Event Studies. In *Prediction Markets*, ed. Leighton Vaughn Williams. Routledge.
- Spears, Brian, Christina LaComb, John Interrante, Janet Barnett and Deniz Senturk-Dogonaksoy. 2009. "Examining Trader Behavior in Idea Markets: An Implementation of GE's Imagination Markets." *The Journal of Prediction Markets* 3(1):17–39.
- Strumpf, Koleman S. 2009. "Introduction to Special Issue on Corporate Applications of Prediction Markets." *Journal of Prediction Markets* 3(1):i–vii.
- Tetlock, Paul. 2004. "How Efficient Are Information Markets? Evidence from an Online Exchange." Columbia University, *mimeo*.
- Tompkins, Robert G., William T. Ziemba and Stewart D. Hodges. 2008. The Favorite-Longshot Bias in S&P 500 and FTSE 100 Index Futures Options: The Return to Bets and the Cost of Insurance. In *Handbook of Sports and Lottery Markets*, ed. Donald B. Hausch and William T. Ziemba. Elsevier: Handbooks in Finance series.
- Tziralis, George and Ilias Tatsiopoulos. 2007. "Prediction Markets: An Extended Literature Review." *The Journal of Prediction Markets* 1(1):75–91.
- Vaughn Williams, Leighton, ed. 2005. *Information Efficiency in Betting Markets*. Cambridge University Press.
- West, Kenneth D. 1996. "Asymptotic Inference About Predictive Ability." *Econometrica* 64(5):1067–1084.
- West, Kenneth D. 2006. Forecast Evaluation. In *Handbook of Economic Forecasting*, ed. Graham Elliott, Clive W.J. Granger and Allan Timmermann. Vol. 1 Elsevier chapter 3, pp. 99–134.
- Wolfers, Justin and Andrew Leigh. 2002. "Three Tools for Forecasting Federal Elections: Lessons From 2001." *Australian Journal of Political Science* 37(2):223–240.
- Wolfers, Justin and Eric Zitzewitz. 2004. "Prediction Markets." *Journal of Economic Perspectives* 18(2):107–126.
- Wolfers, Justin and Eric Zitzewitz. 2006. "Interpreting Prediction Market Prices as Probabilities." NBER Working Paper Series # 12200.
- Wolfers, Justin and Eric Zitzewitz. 2008a. Prediction Markets in Theory and Practice. In *The New Palgrave Dictionary of Economics, Second Edition*, ed. Larry Blume and Steven Durlauf. London: Palgrave.
- Wolfers, Justin and Eric Zitzewitz. 2008b. "Using Markets to Inform Policy: The Case of the Iraq War." *Economica* 76(302):225–250.
- Wright, Jonathan and Jon Faust. 2012. Forecasting Inflation. In *Handbook of Economic Forecasting, Volume 2*, ed. Graham Elliott and Allan Timmermann. Elsevier: Handbooks in Economics series.

Zhou, Guofo and David Rapach. 2012. Forecasting Stock Returns. In *Handbook of Economic Forecasting, Volume 2*, ed. Allan Timmermann and Graham Elliott. Elsevier: Handbooks in Economics series.

Zingales, Luigi. 2010. "Learning to Live with Not-so-Efficient Markets." *Daedalus* 139(4):31–40.

Zitzewitz, Eric. May 30, 2007. "Manipulation can Affect Prices." MidasOracle.org, <http://www.midasoracle.org/2007/05/30/manipulation-can-affect-prices/>.