

NBER WORKING PAPER SERIES

HOW PREDICTION MARKETS CAN SAVE EVENT STUDIES

Erik Snowberg
Justin Wolfers
Eric Zitzewitz

Working Paper 16949
<http://www.nber.org/papers/w16949>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2011

We thank Jim Johnson, Matthew Kahn, Neil Malhotra and Stephanie Rickard for encouragement. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2011 by Erik Snowberg, Justin Wolfers, and Eric Zitzewitz. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

How Prediction Markets Can Save Event Studies
Erik Snowberg, Justin Wolfers, and Eric Zitzewitz
NBER Working Paper No. 16949
April 2011
JEL No. A2,C5,D72,G14,H50

ABSTRACT

This review paper articulates the relationship between prediction market data and event studies, with a special focus on applications in political economy. Event studies have been used to address a variety of political economy questions from the economic effects of party control of government to the importance of complex rules in congressional committees. However, the results of event studies are notoriously sensitive to both choices made by researchers and external events. Specifically, event studies will generally produce different results depending on three interrelated things: which event window is chosen, the prior probability assigned to an event at the beginning of the event window, and the presence or absence of other events during the event window. In this paper we show how each of these may bias the results of event studies, and how prediction markets can mitigate these biases.

Erik Snowberg
Division of the Humanities
and Social Sciences 228-77
California Institute of Technology
Pasadena, CA 91125
snowberg@caltech.edu

Eric Zitzewitz
Dartmouth College
6016 Rockefeller Center
Hanover, NH 03755
Eric.Zitzewitz@Dartmouth.edu

Justin Wolfers
Business and Public Policy Department
Wharton School, University of Pennsylvania
3620 Locust Walk
Room 1456 Steinberg-Deitrich Hall
Philadelphia, PA 19104-6372
and NBER
jwolfers@wharton.upenn.edu

1 Introduction

Event studies have been used in political economy to study the cost of regulation (Schwert, 1981), the value of political connections (Roberts, 1990*a*; Fisman, 2001), the effect of political parties on defense spending (Roberts, 1990*b*), the importance of rules in congressional committees (Gilligan and Krehbiel, 1988), the reaction of different interests to trade legislation (Schnietz, 2000), how party control in parliamentary systems affects broad based stock indices (Herron, 2000), the value of defense contracts (Rogerson, 1989), the effect of the political party of the U.S. President and congressional majorities on particular industry segments (Mattozzi, 2008; Knight, 2006; Herron et al., 1999; Den Hartog and Monroe, 2008; Monroe, 2008; Jayachandran, 2006), and other questions.

Yet event studies are under-used and on the decline. Many of the studies cited above mention the paucity of event studies in political economy, despite their power in revealing the economic impacts of political events. Indeed, due to the issues detailed in this paper, event studies have even fallen out of favor in their traditional strongholds: economics and finance. In what follows we demonstrate, in political contexts, several problems with event studies. We then show that by pairing traditional event studies with prediction markets, to form a *prediction market event study*, many of the problems that have plagued event studies can be mitigated, clearing the way for broader application of event studies in political economy.

At its core, an event study tries to ascertain the economic impacts of a particular event, such as a policy change. These impacts are then used to judge political theories that make different predictions about either the direction or size of an economic effect of a political event, or how an event will differentially affect different economic sectors.

To begin an event study, a researcher first chooses which economic indicator(s) he or she is interested in. If the researcher is interested in the effect of a policy change on a company's performance, he or she will likely examine the price of that company's stock s .¹

¹Event studies and the analysis of prediction markets rely on the efficient markets hypothesis, namely that all available information is quickly reflected in the price of financial securities. For some problems with this hypothesis see Malkiel (2003). Regardless of these potential problems, empirically speaking prediction

The researcher will then select some time t before the policy change occurred, and some time t' after the policy change. The period between t and t' is called the *event window*.

The difference in the price of the stock between the beginning and the end of the event window is proportionate to the impact of the policy change on the price of the company's stock. However, policy changes rarely appear suddenly. Instead, they are created over many months or years. In order to judge the full impact of the policy change on the company's stock, the researcher must judge the *prior probability* of the event at the beginning of the window. We call the chosen prior probability π . Since the posterior probability of the event is always 1, the full effect E of the policy change on the company's stock is given by:²

$$E = \frac{s_{t'} - s_t}{1 - \pi} \quad (1)$$

where s_t is the price of the company's stock at time t .

It should be clear that if the researcher chooses a different start time t , end time t' , or prior probability π , the event study can produce vastly different results in terms of the size and even sign of the estimated effect E . In particular, the longer the event window, the more likely it is that other, unrelated events will occur which will bias the results.

Prediction markets can mitigate these problems. While prediction markets may take many forms, here we focus on the market for a contract that pays \$1 if a certain event, such as a policy change, happens and zero otherwise.³ The price at any given time thus represents the market estimate of the probability of that event happening.⁴

Prediction markets give the researcher an accurate measure of π , the prior probability of an event happening, and may also help to identify an appropriate event window. Moreover, changes in the company's stock price that are unrelated to changes in the probability of the

markets have proven to be extremely accurate (Wolfers and Zitzewitz, 2004; Berg et al., 2008).

²This is a simplified version of the full statistical methodology. For a complete treatment see Schwert (1981).

³For a summary of different types of prediction markets see Wolfers and Zitzewitz (2004).

⁴Wolfers and Zitzewitz (2006) show that under reasonable assumptions of trader risk-attitudes, prediction market prices accurately reflect underlying probabilities.

policy change will show no correlation with changes in the price of the prediction market contract. Thus, the event window does not need to be carefully chosen to exclude other events. Finally, as the probability of an event may change many times in response to political events, each probability change can be analyzed as a separate event where the change in probability of the policy change is accurately measured. By separating a single event window into many small sub-windows and then taking an appropriately weighted average effect across the sub-windows, prediction markets allow for more precise estimates than traditional event studies.

This paper consists of three examples from our previous research (Snowberg, Wolfers and Zitzewitz, 2007*a,b*).⁵ Each example illustrates a particular problem with traditional event studies and shows how the inclusion of prediction markets produces a more accurate estimate of the economic effect of political events. It should be noted that the three issues with traditional event studies are all inter-related—so each example will contain some elements of the other issues. Where possible, we show the difference between our results and research using traditional event studies or misusing prediction markets.

The examples in this paper provide insight into several questions as well as demonstrating the methodological usefulness of prediction markets. First, we show that in the 2004 U.S. Presidential election, candidate convergence did not occur, as predicted by Downs (1957) and many other models. Specifically, the stock market rose 2% in value on news of a Bush victory (over Kerry). Secondly, we show this difference of 2% between Republicans and Democrats has been remarkably consistent over time, appearing in an analysis of all elections between 1880 and 2004. This suggests that whatever the changes in party structure and policy issues over that period, Republicans have consistently been the party of capital, and Democrats the party of labor. Finally, we show that the stock market declined in response to the news of a Democratic victory in the Senate (and House) in 2006, suggesting that,

⁵This research was not the first to use prediction markets to examine the impacts of politics on economic variables. Earlier examples include Slemrod and Greimel (1999), Herron (2000), Leigh, Wolfers and Zitzewitz (2003) and Wolfers and Zitzewitz (2008).

contrary to conventional wisdom, markets do not prefer divided control of the legislature and executive to unified control of both branches.

2 Choosing the Event Window

To illustrate the importance of the event window in traditional event studies, we focus on election night 2004. The price of an S&P 500 future, and the InTrade prediction market tracking Bush's probability of re-election are shown in Figure 1.⁶ Around 3 p.m. EST on election day, the S&P 500 and the probability of Bush's re-election declined in response to leaked exit polls. These polls, which were improperly analyzed (or not analyzed at all) by the news media, showed a Kerry lead in many battleground states. As the actual election results were tallied and Bush's victory became apparent, both the S&P 500 and the probability of Bush's re-election (as revealed by the prediction market) rose.

Combining data from the prediction market and the S&P 500 makes it clear that the S&P 500 rose in response to a Bush victory over Kerry. Moreover, the prediction market allows us to fix an exact scaling between the rise in the probability of a Bush presidency and the rise in the S&P 500. For example, in the first natural experiment in Figure 1 from 3–5 p.m., the probability of a Bush victory decreased by approximately 30% while the S&P 500 decreased by approximately 0.7%. Plugging these values into (1) (appropriately modified so that the denominator reflects the probabilities before and after the event, rather than a value of 1 after the event) yields that the S&P 500 would have declined by 2.3% on news of a Kerry victory (over Bush). Repeating the same exercise for the second natural experiment between 8 p.m. and 1 a.m. the next day, where Bush's chances of re-election rose 65% and the S&P 500 rose 1.3%, implies that the market rose 2.0% on news of a Bush victory (over Kerry). Repeating the same exercise every half hour for the time period in Figure 1 and

⁶At the time of the 2004 election, InTrade was known as TradeSports. We use the company's current name to lessen confusion. As US stock exchanges are closed during the evening and night, we use an S&P future with a near-month delivery date rather than the actual S&P 500.

averaging suggests that the market rose 2.1% on news of a Bush victory (s.e. 0.5%).⁷

Without the information from prediction markets, what might a researcher make of the S&P 500's movement on election night? As previously mentioned, a researcher would have to choose the beginning and end of an event window, as well as the probability that Bush had of winning the Presidency at the beginning of the event window. While it is hard to predict what any given researcher would choose to do in this situation, there are some points that are focal. For example, the opening of the market on election day at 9:30 a.m., the point immediately before the release of any exit polling at 3 p.m. and the close of polls in many east coast states at 8 p.m. are focal.

What about the end of the event window? The most obvious end point is when the major networks called the election for Bush. However, the networks hesitated to call the election in 2004, so a better selection would be when the networks called Ohio for Bush, at 1 a.m. the day after the election.

Two other points are focal end points: the opening of the market the day after the election (9:30 a.m.) and the closing of the market the day after the election (4:00 p.m.—which would give the markets more time to fully respond to the election news).⁸ We will examine the effect of the prior probability of Bush's victory in the next section: for now we set it at 50% as the election was generally considered to be closely contested up until election day.

⁷More technically, we estimated a first differences model of the form:

$$\Delta(\text{S\&P 500})_t = \alpha_t + \beta_t \Delta(\text{Prob. Bush re - election})_t + \varepsilon_t$$

where differences were taken over 30 minute periods. In order to correct for hetero-skedasticity introduced by the occasional missing observation, we weighted by one over the square root of the number of periods since the last observation. We also use White (1980) hetero-skedastic consistent standard errors. The estimates produced via this method for the S&P 500 are roughly consistent with other broad market indices, see (Snowberg, Wolfers and Zitzewitz, 2007a) for details. Readers may be concerned that these results may mask large sector increases or decreases in stock market value. Dooley (2008) shows this was not the case, at least in 2004.

⁸Note that even after the networks had called Ohio for Bush the prediction market hovered between a 90% and 95% chance of a Bush victory, suggesting that the markets might have learned something about the uncertainty of election results from the 2000 recount. A traditional event study could not take account of this nuance.

Note also that the last two possible closes of the event window are especially focal if to researchers using only actual market prices rather than the prices from futures markets.

Using (1), the middle column of Table 1 presents the effect of electing Bush, rather than Kerry, on the S&P 500 for the nine different possible event windows defined by focal start and end points. These nine different event windows provide nine different estimates that range from -1.5% to 2.6%. The standard errors are also much larger than those of the prediction market event study. The highest t-statistic in the table is 1.5, which implies that that estimate is not statistically significant at the 10% level. That is, none of the estimates in Table 1 are significant at conventional levels.

One event window yields the same estimate as the prediction market method (2.1%). While this is heartening, note that this window is just one window out of nine that could be focal, and the fact that this window is correct is somewhat random. Moreover, this is the second highest of the nine possible estimates (in the middle column). Further, three of the estimates in this column are negative—that is, our hypothetical researcher might even get the direction of the effect wrong; that is that markets *rose* on news of a Kerry (rather than Bush) Presidency. Finally, note that the availability of intraday data allows much greater flexibility in the choosing of event windows. This adds researcher discretion and potential bias in traditional event studies, while making prediction market event studies more precise.

3 Prior Probabilities

In the previous section we set the prior probability of Bush's victory at 50%, assuming that our hypothetical researcher thought that the election was a toss-up. Other hypothetical researchers might believe that Bush's probability of winning re-election was closer to 30% if they were a Kerry partisan or as high as 90% if they used the results of Fair (1978, 1996, 2004) as a prediction.

The first and third columns of Table 1 examine the same nine event windows as before for the prior probabilities of 30% AND 90%. Of the 27 different possible answers from a traditional event study, only one gives the correct number, 2.1%. The range of possible

answers is also quite wide, as it may be as low as -7.5% or as high as 12.8%. This twenty percentage-point range is roughly centered on the correct answer, but that doesn't give much hope that a traditional event study give the correct number.

Event study estimates are significantly more noisy without prediction markets to help time and calibrate them. One might wonder, though, whether this noise is sufficiently mitigated by the law of large numbers when one has a large number of similar events and is only interested in the average event effect. For instance, if one looked at a large number of presidential elections and assumed a pre-election probability of 50% in each case, would these errors cancel out?

This is the approach of Santa-Clara and Valkanov (2003), which finds no difference between the stock market's reaction to a Democrat or Republican winning the Presidency. Their methodology is equivalent to conducting an event study for each election between 1928 and 1996, defining the event window as the period between the market close the day before the election and the market close the day after the election.⁹ They chose 50% as the prior probability that the Republican candidate will win each election, and use either 0 or 1 as the posterior probability of the Republican winning (as at the end of the event window the Republican will have either won or lost). They then average the results of these event studies to get a mean and standard deviation, which can be used for statistical tests.

This methodology is also equivalent to regressing the percentage change in the stock market during the event window on the change in the probability of a Republican victory during the event window. This is illustrated in the first panel of Figure 2. This results in a coefficient of 1.3% and a standard error of 0.9%. Thus, the coefficient is not statistically significant at conventional levels.

In the second panel of Figure 2, we instead use prior probabilities gleaned from the historical prediction markets described in Rhode and Strumpf (2004, 2006, 2008). These markets, run on the curb exchange on Wall Street, were the dominant form of election

⁹Until 1984 the market was not open on election day.

projection before the advent of scientific polling. We use these markets to establish the true prior probability of a Republican winning the presidency each election. In this panel, the relationship between the electoral fortunes of Republicans and the stock market is clearly much stronger. By using prediction market event studies, we find that the markets rose 3.0% in response to a Republican victory between 1928 and 1996. The standard error of this estimate is 1.2%, making the coefficient statistically significant at the 5% level.

Why are our results so different from those in Santa-Clara and Valkanov (2003)? Using a 50% prior probability of Republican victory for each election, even if correct on average, adds measurement error to the right-hand side of the regressions. This biases coefficients towards 0. By using an accurate estimate of the prior probabilities of Republican victory gleaned from prediction markets, we are able to recover the correct coefficient, which is similar to the result from 2004.

4 Parceling out Market Movements Related to Politics

Even if a researcher picked the correct event window and prior probability, the results of a traditional event study could still be biased by market movements caused by other events. While financial theory says that, on average, market movements due to other events should have very little effect, the probability that they actually have no effect over a moderate length of time (a day) is quite small.

To illustrate this, consider the 2006 U.S. Congressional election, as illustrated in Figure 3. The first panel displays the entire course of election night, showing the movements of a prediction market contract tied to whether Republicans maintained a majority in the Senate, another contract tied to whether the Republicans maintained a majority in the House, and an S&P 500 near-month future.

At the beginning of election night, Republicans had less than a 20% chance of maintaining their majority in the House. This slowly converged to zero as the results of the few close

contests became known. The Republicans in the Senate seemed to be in a better position to hold on to their majority, entering the evening with a 70% chance that rose above 90% when early exit polls favored them. However, when the vote totals of the last counties came in for Virginia and Missouri, the markets quickly reversed themselves, sending the probability of a Republican Senate majority to 10%.

A cursory glance at these three financial contracts together would suggest that there was little relationship between Republican (vs. Democratic) majorities in the House and Senate and the S&P 500. The markets did trend downwards over the entire evening which may have led a traditional event study (even with the proper prior probabilities) to conclude that the change of the House from a Republican to Democratic majority would result in a *decrease* in the S&P 500 of over 2%. However, if the event window included Rumsfeld's resignation, a traditional event study would have concluded that the change of the House from a Republican to a Democratic majority resulted in more than a 2% *increase* in the S&P 500.¹⁰

The truth, as revealed by prediction markets, is strikingly different. The second panel of Figure 3 shows a close-up of the time period when there was a large change in the probability of a Republican majority in the Senate. This panel shows that the market moves in lock-step with the probability of a Republican majority in the Senate, although the movement is small. Moreover, the market's non-response to the swing of Virginia's Senate seat into the Democratic column reveals that the movement in the S&P 500 is due to a change in the majority status of the parties, not a shift in the median voter of the Senate.

By employing the same methodology as that used in our study of the 2004 election we find that a shift in the Senate from a Republican to a Democratic majority caused the S&P 500 to decline by 0.2% points (standard error 0.09%—making the result statistically significant at the 10% level), while the S&P 500 was statistically indifferent to the change

¹⁰InTrade ran a contract on the probability of Rumsfeld's resignation. Although thinly traded, it suggests that his resignation lead to a 1.5% increase in the S&P 500. It should be noted that this rise may not have been caused entirely by Rumsfeld's resignation, but may be a response to some sign of a change in the Bush administration's policy in Iraq and the Middle East more broadly.

in the House majority.¹¹ It should be noted that these results are from a regression that includes *the entire* time covered by the first panel of Figure 3, further underscoring the power of prediction markets to parcel out the market response to election news from changes in the index due to other factors.¹² Moreover, this parcelling reduces noise, allowing more precise estimates. To see this, compare the standard errors in Table 1 to the standard error of 0.5% of the prediction market event study over the same period.

5 Misuse of Prediction Markets

Each prediction market contact is a new variable. A new variable will have correlations with other variables, and thus chances to mistake correlation for causation. To date, the most common instance of such confusion involving prediction markets comes from studies that regress stocks and stock market indices on prediction markets in the weeks and months before an election.

Figure 4 presents a stylized, but relatively complete, representation of the system relating economic and political events, voters, prediction markets and financial markets. Researchers have generally been interested (as we are in this paper) in discovering the effect of the electoral fortunes of various candidates on financial markets. This effect is represented by the large arrow in the figure.

Estimating the effect of politicians on the market is complicated by the fact that the health of the economy in general, and financial markets in particular, have a profound impact on the (re-)election chances of politicians.¹³ A simple regression of the stock market on prediction market prices will be biased by reverse causality. For example, Herron (2000),

¹¹A similar study in 2002 showed that the S&P 500 increased by 0.6% in response to a shift from a Democratic to a Republican Senate. See Snowberg, Wolfers and Zitzewitz (2007*b*).

¹²Moreover, the results do not change if the contracts for the House and Senate are entered together or separately, indicating that this is not the result of a joint effect of the loss of *both* houses of Congress.

¹³The academic literature on economic voting Kramer (1971); Fiorina (1981); Lewis-Beck (1988) is truly staggering. See (Lewis-Beck and Paldam, 2000) for a review. Particularly relevant here is the theory of sociotropic (Kinder and Kiewiet, 1979, 1981) voting which is based on the observation that the perception of general economic variables such as the stock market is strongly related to vote choice, whereas personal economic experience is less strongly correlated.

which studies the relationship between Labour's fortunes in the 1992 British parliamentary elections and the FTSE 100, the standard British stock index, speaks about this issue, but is forced to assume that reverse causality is not an issue due to data limitations. Reverse causality is likely here, so the results of this study are likely biased.¹⁴

A straightforward illustration of this bias is found in Figure 5 which shows the S&P 500 and Bush's probability of re-election over the course of the 2004 election. Note that the variables are initially uncorrelated, and then become correlated as the election nears its conclusion. This mirrors the finding in Fair (1996), that economic performance in the two quarters preceding a presidential election are important predictors of presidential vote shares.

Regressing the S&P 500 on the probability of Bush's re-election over the entire time frame of Figure 5 yields estimates of Bush's impact on the S&P 500 (vs. Kerry) that are as high as 24%. This stands in stark contrast with the actual effect of Bush on the S&P 500 of only 2%.

Recognizing the possibility for such bias studies such as Herron et al. (1999), Knight (2006), Mattozzi (2008) instead regress industry stock portfolios on prediction market prices. It is unlikely that the economic performance of a single industry would have a significant impact on the probability of Bush or Kerry's election, side-stepping concerns about reverse causality. However, such studies may present other biases, also illustrated in Figure 4.

Specifically, Figure 4 shows that other factors may affect both an industry's performance and a candidate's fortunes. This is best illustrated with an example taken from Snowberg, Wolfers and Zitzewitz (2007*a*). Suppose that an election features a pro- and anti-war candidate, and the pro-war candidate is a more capable war president. If shares in defense contractors increase in value when the pro-war candidate's electoral prospects improve, one might be tempted to conclude that the defense contractors stocks are worth more because

¹⁴Note that classical event studies will not always be biased by reverse causality. For example, in Slemrod and Greimel (1999) it is unlikely that the municipal bond market had any effect on the outcome of the Republican presidential primary, and in Wolfers and Zitzewitz (2008) it is unlikely that the price of oil changed the probability that the U.S. invaded Iraq.

there is a higher chance of the pro-war candidate will be elected. However, a third factor—such as threatening actions from another nation—may have led both numbers to appreciate: the defense contractor’s from their increased sales in an increasingly likely war, and the pro-war candidate’s from his country’s increased need of his leadership in wartime.

Such pre-election analyses are popular with political commentators and the press. Indeed, in July 2008, a month after Obama had clinched the Democratic nomination (and four months before the general election) commentators noted the correlation between Obama’s increasing chance of winning the presidential election and the declining stock market illustrated in Figure 6. The correlation between the data series in Figure 6 is slightly more than -0.9. Some commentators even went so far as to describe this pattern of equity prices as the “Obama Slump.”

Given the evidence in this paper and our previous research, It is far more likely that a declining economy was hurting both the stock market and McCain’s chance of election. This did not prevent outlets such as the *Wall Street Journal* from publishing such speculation on their website.¹⁵

As with all endogeneity concerns, the remedy is some form of instrumental variable or a natural experiment. While it is theoretically possible to construct a prediction market that would be correlated with one variable and uncorrelated with another in the analysis, and thus be able to serve as an instrumental variable (Wolfers and Zitzewitz, 2005), to date we are unaware of a successful application of this procedure. Instead, we focus on *instrumental events* such as debates or gaffes (like the infamous “Dean Scream”) that affect the political fortunes of the candidates without directly effecting the economy. By focusing on election night in this paper, we examine a period when votes have already been cast, and thus the information that is revealed as votes are counted can only affect the economy through changes in expectations, not through changes in vote choice as illustrated in Figure 7.

¹⁵See <http://online.wsj.com/article/SB121592969771748931.html> and <http://www.midasoracle.org/2008/07/15/intradew-jones/>.

6 Conclusion

While we have documented problems that plague many event studies, we do not mean to imply that all event studies without prediction markets will be flawed. For example, reverse causality may not be a problem when the economic object of study is unimportant for electoral outcomes, such as in Slemrod and Greimel (1999), which studies municipal bond markets. To take another recent example, Jayachandran (2006) studies the market response to different majority parties in the Senate by using an instrumental event, the switch of Senator Jeffords (I-VT) from the Republican party to an independent, which handed control of the Senate to the Democrats. To conduct her study, she created portfolios of stocks associated with Democrats and Republicans, and then observed changes in market value as news about Jeffords' switch leaked out over the span of a week. Her study shows clearly that companies associated with Democrats increased in value relative to those associated with Republicans. What she could not study was what happened to the market as a whole as many other events would have increased or decreased the value of a broad market index like the S&P 500. So while Democratic stocks increased slightly in value relative to Republican stocks, her study leaves open the possibility that all stocks rose (or fell) substantially due to Jeffords' switch.

Prediction markets erase this constraint. As we have shown in this paper, broad market indices move very little in response to a switch in the majority party of the Senate. By augmenting event studies with prediction markets other scholars will no doubt come up with creative ways to address many other unanswered questions.

References

- 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.
- Den Hartog, Chris and Nathan W. Monroe. 2008. "The Value of Majority Status: The Effect of Jeffords' Switch on Asset Prices of Republican and Democratic Firms." *Legislative Studies Quarterly* 33(1):63–84.
- Dooley, Patrick. 2008. "Elections and Market Capitalization Effects: Evidence from the 2004 US Presidential Election." Dartmouth College, *mimeo*.
- Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York, NY: Harper Collins.
- Fair, Ray C. 1978. "The Effect of Economic Events on Votes for President." *The Review of Economics and Statistics* 60(2):159–173.
- Fair, Ray C. 1996. "Econometrics and Presidential Elections." *The Journal of Economic Perspectives* 10(3):89–102.
- Fair, Ray C. 2004. "Presidential Vote Equation—October 29, 2004." Yale University, *mimeo*, <http://fairmodel.econ.yale.edu/vote2004/vot1004.htm>.
- Fiorina, Morris P. 1981. *Retrospective Voting in American National Elections*. New Haven, CT: Yale University Press.
- Fisman, Ray. 2001. "Estimating the Value of Political Connections." *The American Economic Review* 91(4):1095–1102.
- Gilligan, Thomas W. and Keith Krehbiel. 1988. "Complex Rules and Congressional Outcomes: An Event Study of Energy Tax Legislation." *Journal of Politics* 50(3):625–654.
- Herron, Michael C. 2000. "Estimating the Economic Impact of Political Party Competition in the 1992 British Election." *American Journal of Political Science* 44(2):326–337.
- Herron, Michael C., James Lavin, Donald Cram and Jay Silver. 1999. "Measurement of Political Effects in the United States Economy: A Study of the 1992 Presidential Election." *Economics and Politics* 11(1):51–81.
- Jayachandran, Seema. 2006. "The Jeffords Effect." *The Journal of Law and Economics* 49(2):397–425.
- Kinder, Donald R. and D. Roderick Kiewiet. 1979. "Economic Discontent and Political Behavior: The Role of Personal Grievances and Collective Economic Judgments in Congressional Voting." *American Journal of Political Science* 23(3):495–527.
- Kinder, Donald R. and D. Roderick Kiewiet. 1981. "Sociotropic Politics: The American Case." *British Journal of Political Science* 11(2):129–161.

- Knight, Brian. 2006. "Are Policy Platforms Capitalized into Equity Prices? Evidence from the Bush/Gore 2000 Presidential Election." *Journal of Public Economics* 90(4-5):751–773.
- Kramer, Gerald H. 1971. "Short-Term Fluctuations in US Voting Behavior, 1896-1964." *The American Political Science Review* 65(1):131–143.
- Leigh, Andrew, Justin Wolfers and Eric Zitzewitz. 2003. "What do Financial Markets Think of the War with Iraq?" *NBER Working Paper #9587*.
- Lewis-Beck, Michael S. 1988. *Economics and Elections: The Major Western Democracies*. Ann Arbor: The University of Michigan Press.
- Lewis-Beck, Michael S. and Martin Paldam. 2000. "Economic Voting: An Introduction." *Electoral Studies* 19(2-3):113–121.
- Malkiel, Burton G. 2003. "The Efficient Market Hypothesis and Its Critics." *Journal of Economic Perspectives* 17(1):59–82.
- Mattozzi, Andrea. 2008. "Can We Insure Against Political Uncertainty? Evidence from the U.S. Stock Market." *Public Choice* 137(1–2):43–55.
- Monroe, Nathan W. 2008. "Majority Party Dividends: The Effect of the 1994 Election on Asset Prices of Republican Constituent Firms." Michigan State University, *mimeo*.
- 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. 2006. "Manipulating Political Stock Markets: A Field Experiment and a Century of Observational Data." University of Arizona, *mimeo*.
- Rhode, Paul W. and Koleman S. Strumpf. 2008. "Historical Election Betting Markets: An International Perspective." University of Arizona, *mimeo*.
- Roberts, Brian E. 1990a. "A Dead Senator Tells No Lies: Seniority and the Distribution of Federal Benefits." *American Journal of Political Science* 34(1):31–58.
- Roberts, Brian E. 1990b. "Political Institutions, Policy Expectations, and the 1980 Election: A Financial Market Perspective." *American Journal of Political Science* 34(2):289–310.
- Rogerson, William P. 1989. "Profit Regulation of Defense Contractors and Prizes for Innovation." *The Journal of Political Economy* 97(6):1284–1305.
- Santa-Clara, Pedro and Rossen Valkanov. 2003. "The Presidential Puzzle: Political Cycles and the Stock Market." *The Journal of Finance* 58(5):1841–1872.
- Schnietz, Karen E. 2000. "The Institutional Foundation of US Trade Policy: Revisiting Explanations for the 1934 Reciprocal Trade Agreements Act." *Journal of Policy History* 12(4):417–444.

- Schwert, G. William. 1981. "Using Financial Data to Measure Effects of Regulation." *Journal of Law and Economics* 24(1):121–158.
- Slemrod, Joel and Timothy Greimel. 1999. "Did Steve Forbes Scare the US Municipal Bond Market?" *Journal of Public Economics* 74(1):81–96.
- 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. "Party Influence in Congress and the Economy." *Quarterly Journal of Political Science* 2(3):277–286.
- White, Halbert. 1980. "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity." *Econometrica* 48(4):817–838.
- Wolfers, Justin and Eric Zitzewitz. 2004. "Prediction Markets." *Journal of Economic Perspectives* 18(2):107–126.
- Wolfers, Justin and Eric Zitzewitz. 2005. "Five Open Questions About Prediction Markets". 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 Press.
- Wolfers, Justin and Eric Zitzewitz. 2006. "Interpreting Prediction Market Prices as Probabilities." *NBER Working Paper #12200* .
- Wolfers, Justin and Eric Zitzewitz. 2008. "Using Markets to Inform Policy: The Case of the Iraq War." *Economica - Forthcoming* 76(302):225–250.

Table 1: Different event windows and prior probabilities of Bush's re-election give widely different estimates of the market's response to Bush's victory over Kerry.

Event Window		Prior Probability					
Start (11/2)	End (11/3)	30%		50%		90%	
9:30 a.m.	1 a.m.	0.6%	(.9%)	0.8%	(1.3%)	3.9%	(6.3%)
	9:30 a.m.	0.9%	(1.0%)	1.2%	(1.4%)	6.1%	(6.9%)
	4 p.m.	-1.1%	(.8%)	-1.5%	(1.2%)	-7.5%	(5.9%)
2 p.m.	1 a.m.	1.0%	(.7%)	1.4%	(1.0%)	7.0%	(4.8%)
	9:30 a.m.	1.3%	(.9%)	1.8%	(1.2%)	9.2%	(6.0%)
	4 p.m.	-0.6%	(.5%)	-0.9%	(.7%)	-4.4%	(3.6%)
8 p.m.	1 a.m.	1.5%	(1.2%)	2.1%	(1.6%)	10.6%	(8.2%)
	9:30 a.m.	1.8%	(1.2%)	2.6%	(1.7%)	12.8%	(8.5%)
	4 p.m.	-0.1%	(1.0%)	-0.2%	(1.4%)	-0.8%	(6.8%)

Notes: Standard errors (in parenthesis) are the standard deviation of price movements in the near quarter CME S&P 500 future for 60 days before and after November 2, 2004, appropriately adjusted by the prior probability.

Figure 1: Prediction markets reveal that Bush's re-election on November 2, 2004 increased equity prices by approximately 2%.

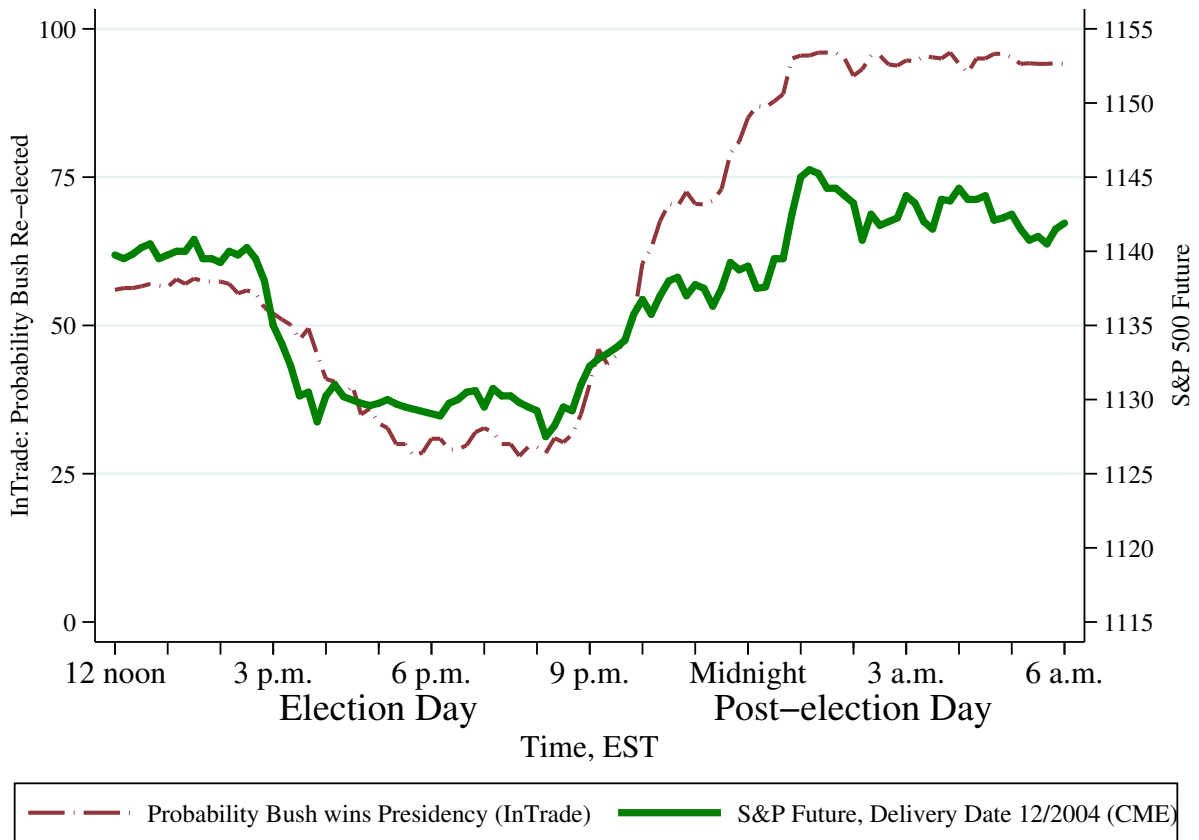


Figure 2: The long-term relationship between elections and equities is obscured without prediction markets.




Figure 3: Prediction markets separate relevant equity movements from irrelevant movements.




Figure 4: The interactions between information, voters, prediction markets and financial markets are complicated.





Figure 5: Bush's chances of re-election and the S&P 500 become increasingly correlated over time.



Notes: Graph displays a five-day moving average of the closing prices at 4 p.m. EST for both the S&P 500, and the InTrade prediction market contract on Bush's re-election.

Figure 6: The high negative correlation between equities and Obama’s chance of winning are likely due to poor economic performance hurting McCain.




Figure 7: On election night, information can only effect economic variables through the expectations about each candidate’s chances of election.

