Understanding Elections:

Measuring Electoral Determinants with Electronic Prediction Markets

S. Travis May

Presented to the Department of Economics in partial fulfillment of the requirements for a Bachelor of Arts degree with Honors

Harvard College

Cambridge, Massachusetts

March 19, 2009

Abstract

While significant research has demonstrated that prediction markets are highly accurate indicators of the real-time status of an election, few studies within political economics or political science have used them as a tool to understand voting behavior. Assuming their accuracy, we utilize these markets to measure the primary determinants of the 2008 United States election. First, we use electronic prediction markets to show that the financial crisis had a decisive impact on the outcome of the presidential election, and conversely that Obama's victory had a sizable impact on equity markets. Next, we use prediction markets to measure the scope of the impact of the presidential campaign. We find that campaign events had a significant – but small – impact on state-by-state outcomes, but campaign spending had no significant local-level effects. Finally, we use data from electronic prediction markets to measure the size of presidential "coattail effects" on congressional elections, finding strong effects in House elections but weak effects in the Senate. Throughout each of these chapters, our use of data from electronic prediction markets enables us to circumvent the typical measurement problems for each of these questions by increasing the sample size and eliminating endogeneity biases with a real-time metric of election status.

Acknowledgements

This paper would not have been possible without the invaluable ideas and support of Yiling Chen, whose expertise with prediction markets and feedback were instrumental throughout the writing process. I am thankful to Dan Wood, whose extensive advice and comments were formative in shaping my paper from the beginning. I am also appreciative of the ideas I gained from conversations with James Stock and Andrei Shleifer.

Additionally, I am grateful to Intrade for its generous sponsorship of electronic prediction market data and DiskTrading.com for providing financial futures market data. Publically-available data from Iowa Electronic Markets were also instrumental in performing this analysis.

Finally, I am extremely grateful for the stylistic support provided by Greg, Eva, and Holly, who were extremely generous with their last-minute editing and feedback.

Table of Contents:

Chapter I

Introduction

"Experimental prediction markets were established at the University of Iowa in 1988, and they have since amassed a very impressive record, repeatedly outperforming the polls. Research by economic historians has documented betting on elections over a century ago, and the impressive forecasting record of prediction markets was also evident in the period before scientific polling was adopted…

Through the process of different people trading based on their own observations about the race, prediction markets prices come to aggregate disparate pieces of information into a single summary measure of the likelihood of various outcomes. Moreover, if this market operates efficiently, it will appropriately summarize all of this information and the price will become the most statistically accurate forecast of the election outcome."

--The Wall Street Journal

As the cornerstone of democracy, elections play a critical role in the United States political system. Every four years, star candidates rise and fall, scandals emerge and disappear, and long campaigns are fiercely fought before settling on a winner. Pundits devise dozens of stories that "explain" what determined the winner, and television and radio talk shows will debate these theories for months. But to the academic empiricist, understanding an election is problematic. To understand the determinants of a *specific* election, the sample size is only one event – so how can we know whether John McCain would have fared better with a different vice-presidential candidate? To understand election determinants, political economists like Ray Fair have pooled the data of many elections and constructed models of very broad trends; yet because of the small sample size and the complexity of the variables impacting elections, these models yield only a few answers about what shapes elections in general, and much less about what shapes specific elections.

In this paper, we propose the use of a new tool for understanding elections: electronic prediction markets. These markets have been studied at length, with a broad theoretical and empirical consensus that they yield reasonably unbiased, efficient estimators of the probability that an event occurs. Instead of focusing on their efficiency, as most papers have, we will instead *assume* that they are efficient. In this paper, we will attempt to answer the following question: given that electronic prediction markets are efficient, what can they tell us about the 2008 election? We will explore several determinants of election success, each of which is an area that has proven difficult to measure without the use of prediction markets. While these chapters each yield interesting results, our focus is primarily methodological in understanding how the

existence of a continuous metric of election probabilities can be useful for causal analysis.

Background on Electronic Prediction Markets

Electronic prediction markets are futures exchanges in which the value of an asset is tied to the outcome of a particular event. In a vote share market, traders bid in a continuous double auction for the optimal price that they will pay for a contract that pays $1 for every percentage point that a particular party receives. For example, in the 2000 election, George Bush received 47.9% of the vote, and thus the Republican contract paid off $47.90 at the end of the election. In a winner-take-all (WTA) election, traders bid on contracts for which candidate will win the election and receive $100 per share of the winning candidate held.

In a system with no transaction costs and risk neutrality, Wolfers and Zitzwitz (2004) show that the market price of the contract should equal the median market participant's expected election outcome weighted by trading volume. Furthermore, Berg et. al. (2001) shows that the market price is a remarkably accurate short-term predictor of the real election outcome, consistently outperforming polls as a last-day prediction tool. In a follow-up study, Berg et. al. (2003) finds that electronic markets have exceptional long-term predictive abilities, and greatly outperform both polls and forecasting models.

The intuition behind these results is clear: traders are forced to "put their money where their mouth is," and they are thus incentivized to be accurate. Since projection models and poll results are generally publically available, that information can be incorporated into the market price. Thus the market prediction should be at least as

accurate as any type of publicly-known prediction mechanism, if not more so. Interestingly, this assumption does not depend on traders being a random sample of the population, as the market can incorporate information without all voters participating. As Berg et. al. (2001) points out, market participants are far from a random sample; they tend to be wealthy, young, and highly-educated. Furthermore, the average trader in the market can be biased and trade based on his own bias; empirically, there is a small set of "marginal traders" who are the ultimate price-makers and hunt for arbitrage opportunities – and their probabilistic assumptions are quite accurate.

With this consensus of research demonstrating that prediction markets are accurate estimators, we intend to utilize the market results in this paper as a continuous set of election results – changing over the course of the election only as the expected outcome changes. Thus, we have a much more complete data set, with a continuous set of data points that can be evaluated in order to measure the importance of contextual events.

Determinants of Election Results

In the chapters ahead, we will make use of electronic prediction markets to explore several determinants of election results. Each chapter explores a question about the 2008 election that existing methods have been unable to capture fully.

In Chapter II, we develop a simple model to help explain election phenomena. In the model, the median voter maximizes her utility in a two-candidate election by selecting whether to vote and for whom. Throughout the paper, we will explore the determinants of her vote in more detail.

In Chapter III, we examine the role of the financial crisis and economic phenomena in the 2008 U.S. presidential election. Much research has focused on the impact of economics on elections in general, but the degree to which economics matters may change from election to election. While Fair's model has shown the impact of the economy *in general*, electronic prediction markets provide sufficient data from a single election to measure the impact the economy had on a *specific* election. By using electronic prediction markets, we measure both the impact of the election on financial markets, as well as the impact of financial markets on the election. We find that the election was in fact largely shaped by financial changes, and that Obama's ultimate victory had a significant impact on markets.

While Chapter III focuses on phenomena that are outside the candidates' control, we turn to the candidates' decisions in Chapter IV. By controlling for the initial probability of winning using electronic prediction market prices, we measure the impact of campaign spending and campaign events on the election. We find a significant – but small – effect of campaign events on the election outcome, but no significant local effect of campaign spending. In this case, electronic prediction markets help solve an endogeneity bias that has impacted other analyses.

In Chapter V, we measure the size of "coattail effects" in the 2008 election, or the impact that the presidential candidates had on congressional candidates. Traditional methods of measuring this value have assumed that all correlation between candidates is due to coattail effects, ignoring other changes that impact both elections such as economic context. By observing changes in electronic prediction markets and using events that only directly affect the presidential election as an instrument, we eliminate

this endogeneity and more accurately measure the scope of coattail effects. We find that coattail effects were large in House elections and small in Senate elections during the 2008 U.S. election.

In each of these studies, the assumption that electronic prediction markets are efficient is the foundation of our analysis. This assumption allows us to boost our sample size and correct biases by tracking a continuous indicator of the election status, presenting a new tool that allows us to better understand our political system.

Chapter II

A Model of Voter Preferences

In this chapter, we will develop a simple model that informs our understanding of election effects throughout this paper. We model voter preferences by examining a two-candidate election. We select two candidates for simplicity in order to rule out strategic voting behavior, but the model can be expanded to n candidates. In a two-candidate election, voter i makes a discrete choice to vote for the candidate that provides the most expected utility upon victory. We create a discrete variable $X_{i,j}$ that is equal to 1 when voter i decides to vote for candidate j, 0 if the voter decides not to vote, and -1 if the voter votes against candidate j. Thus, in deciding whether to vote for candidate j over candidate k in the two-party election, the voter optimizes:

$X_{ij} = 1$ if $U_{dif} \geq \overline{U}$ **(Equation 2.1)**

$X_{ij} = 0$ if $-\overline{U} < U_{dif} < \overline{U}$

$X_{ij} = -1$ if $U_{dif} \leq -\overline{U}$

Where $U_{dif} = E[U_{ij}]-E[U_{ik}]$

X_{ij}: vote cast by voter i for candidate j. Not voting counts as 0 votes, and voting against counts as a negative vote for the candidate
U_{ij}: utility voter i receives in the state of the world where candidate j wins
U_{ik}: utility voter i receives in the state of the world where candidate k wins
$\overline{U}$: absolute utility difference threshold that causes voter i to cast a vote[1]

[1] Much has been written about the "paradox of voting," questioning why voters participate in an election if the probability of casting the winning ballot is negligible. For our purposes, we assume that – for whatever

U_{ij} is equal to the utility derived by voter i if candidate j wins. The rational voter selects the candidate whose policies agree most with her own views, weighted by the importance of each particular view. For example, some individuals are fairly indifferent about a candidate's view on Israeli politics, while for others it is a defining weight. Thus we model U_{ij} as a function of the extent to which the voter agrees with the candidate on policies (v_{ij}) and the weight of importance she assigns to each view (i_i). Thus:

$$U_{ij} = F(i_i, v_{ij}) \qquad \textbf{(Equation 2.2)}$$

$\partial F/\partial v_{ij} \geq 0$
U_{ij}: utility voter i receives in the state of the world where candidate j wins
v_{ij}: vector of degree of agreement between voter i and candidate j's decisions
i_i: vector of weights of importance that a voter gives to particular issues

A voter does not have *a priori* guarantees on how much she will agree with a particular candidate, since candidate positions may change over time and new issues will likely arise that are not directly confronted on the campaign trail. Thus v_{ij} is not directly measurable by a voter, and the voter must make an expectation of this variable in their decision. The factors entering the voter's expectation of v_{ij} include:

(1) election context (c), including factors such as economic conditions that are likely to affect the individual voter's views;

(2) the extent that the voter agrees with the revealed and known policy preferences of the candidate (r_{ij}), which reflect the decisions that the candidate would likely make in office;

---

reasons a voter decides to vote - $\overline{U}$ represents the utility difference threshold needed to vote. This may be exceptionally large for the voter who only cares if they cast the deciding vote, or it may be smaller for citizens with other motives and incentives.

(3) the candidate's personal characteristics (p_j), as a candidate that is considered trustworthy or intelligent may be expected to make more agreeable policy decisions;

(4) other signals about the candidate's future decisions (s_{ij}), such as party affiliation.

At the same time, a voter's decision is impacted by changes in the importance she assigns each topic in the election (i_i), which is a function of personal tastes (t_i) and election context (c). Thus we restate the voter's decision as the following system:

$$E(U_{ij}) = E(F(i_i, v_{ij}))$$ **(Equation 2.3)**
$$i_i = i(c, t_i)$$
$$E(v_{ij}) = v(c, r_{ij}, p_j, s_{ij})$$

U_{ij}: utility voter i receives in the state of the world where candidate j wins
v_{ij}: vector of degree of agreement between voter i and candidate j's decisions
i_i: vector of weights of importance that a voter gives to particular issues
c: context of election
t_i: voter i's tastes
p_j: personal characteristics of candidate j
r_{ij}: historical correlation between voter i and candidate j's known beliefs
s_{ij}: signals of future agreement between voter i and candidate j

Equation 2.3 will be our core model which we will explore and expand upon throughout this paper.

Chapter III

The Economy and the 2008 Election

In this chapter, we examine the impact of economic and financial context on the 2008 U.S. presidential election. During the election, the rapid deterioration of the economy and financial markets altered voter preferences and priorities, impacting election results. At the same time, changes in the election affected financial markets by shaping expectations of future policy changes. To quantify each of these effects, we regress electronic prediction market prices on a variety of economic and financial indicators. In order to isolate the impact of political changes on financial markets, we restrict our analysis to Election Day changes, in which voter decisions are already fixed and both financial and prediction markets respond to incoming election news. We find that a 1% increase in the probability Obama wins the election is associated with a 0.16% increase in S&P 500 futures on Election Day. Applying this result retroactively, we then identify the impact of economic and financial changes on the election. We find that Obama's probability of victory was largely determined by financial changes, as it was heavily boosted by decreases in equity prices and increases in treasury and oil prices. We discuss these results in light of our simple voting model.

Introduction

In September 2008, the popular belief was that John McCain had gained a slight edge in his electoral contest against Barack Obama. A coincidence of factors, ranging from a successful strategy of political aggressiveness to a vice-presidential selection that inspired the conservative base, led to the conclusion that he was the favorite in the race. Yet, by mid-October, analysts had all but called the election in favor of Barack Obama.

What happened to change popular opinion in that short time period? A financial crisis, coupled with an already unpopular President, crippled the incumbent Republican Party, and likely sealed John McCain's electoral fate. Economic distress, an exogenous factor that neither candidate controlled, ultimately shaped the election more than the candidates' personalities or policy positions. In fact, as Fair (1978) demonstrates, this is the norm: elections are largely determined by contextual events that don't necessarily bear on the candidates or their proposed policies.

In Equation 2.3, which shows the expected utility of either candidate winning the election, economic context enters the equation at two points: first, it affects a voter's relative importance weightings between issues (i_i); and second, it alters political views directly (v_{ij}). In an election decision, rational voters may choose to evaluate the candidates based on their parties' historical abilities to create a prosperous economy under their incumbency. Lewis-Beck (2000) proposes the following type of scenario to consider economic effects on each party. Suppose that the Republican Party prioritizes low inflation. Then, in times of high inflation, voters may choose to turn to the Republicans to help keep inflation down. This is called the *clientele hypothesis*, and supposes that each party has particular values that it prioritizes successfully. Yet,

consider the case that the Republicans are incumbents, and despite their typical prioritization of anti-inflation measures, inflation continues to be high. In this case, voters may choose to penalize the party especially hard (and deviate from its ideology), as they failed in keeping their mandate. This is the *saliency hypothesis*. While the clientele hypothesis shows how the economy can cause a reordering of preferences and shift in i_i, the saliency hypothesis establishes a change in political views as the economy changes (v_{ij}).

Beginning with a landmark study by Ray Fair (1978), political economists have sought to measure contextual effects on elections through election forecasting, an empirical methodology that seeks to answer "how, if at all, economic events effect voting behavior." The underlying assumption is that individuals vote at least in part with their purses, and they evaluate "the past economic performances of the competing parties and [vote] for the party that provides the highest expected future utility."

As demonstrated by Fair (1978), forecasting models can measure the importance of various economic effects by regressing the vote share of each party against a wide range of contextual variables. Fair finds that "economic events as measured by the change in real economic activity in the year of the election do appear to have an important effect on votes for president." With a limited number of variables, including factors such as incumbency and recent real GDP growth per capita, forecasting models such as the one in Fair (1996) show a standard error of only 2% in predicting popular vote outcomes.

While these results indicate an overwhelming trend that economic context plays an important role in elections, the conclusions that can be drawn are thin in significance.

After every election, a debate breaks out on how to best improve the models, and why any errors may have occurred – leading to the creation of a refined model. For example, Erikson, Bafumi, and Wilson (2001) show that, even in retrospect, the 2000 election results defy nearly all reasonable models. Under the forecasting paradigm, specific effects within a single election cannot be measured, so there is no way to gauge why a particular election deviated from expectations; in some elections, it is possible that the economy matters more than other elections, and there is no way to assess these types of effects. Furthermore, in general, there is no way to gauge causality, as there are close relationships between many of the regressors involved and outside forces that could shape any of them. The fundamental issue for why more variables are not able to be tested is one of sample size: there is only one election every four years from which to draw data, and there are substantial changes that occur over that time frame. As Fair (1996) describes:

> *"Data mining is a potentially serious problem in the present context, given the small number of observations. Much searching was done in arriving at the final specification, and it may be that an equation was found that fits the historical data well, but that is, in fact, a poor approximation of the way that voters actually behave. Put another way, the equation may be overparameterized: since there are a relatively high number of parameters for the number of observations, small changes in the data or the specification can lead to substantial changes in the estimates."*

Thus, while forecasting models have become an effective tool for predicting election outcomes, they are ineffective at indicating causal effects and are heavily limited by the size of their dataset.

To overcome the limited datasets that forecasting models can currently access, we propose combining the election forecasting methodology with electronic prediction markets. By using the market price as a real-time indicator of an election's status, we vastly expand the available dataset, and we are able to gauge the impact of exogenous events as they occur. In this chapter, we use a prediction market dataset to measure the impact of contextual events on elections, as well as the impact of elections on financial markets.

Our results indicate that the financial and economic events were indeed a decisive factor in the election. On Election Day, we find that incoming news about Obama's success in the election had a strong positive effect on S&P 500 futures. By applying these results backwards, we are then able to identify the effect of financial markets on the election. We find an effect in the opposite direction: increases in Obama's election probability are strongly associated with equity market *decreases*, 5-year treasury price *increases*, and oil price *increases*. In just the period from September 5 to October 15, we predict a 15.9% boost in Obama's probability due to the financial crisis, compared to an actual increase of 19.8% over this time period.

This chapter proceeds as follows: first, we outline an empirical methodology using prediction markets to measure effects of the election on financial markets; next, we reverse this analysis to measure the impact of financial markets on elections; finally, we discuss these results and their implications for election outcomes.

Part I: The Election's Impact on Financial Markets

Despite the intuitive relationship between financial markets and the election, the specifics of this relationship have been elusive to measure. One core problem is the bi-directionality of these effects. Political decisions affect the economy, and financial markets must price in expected policy changes as the circumstances of an election change; simultaneously, the financial markets and economic context can shape election outcomes.

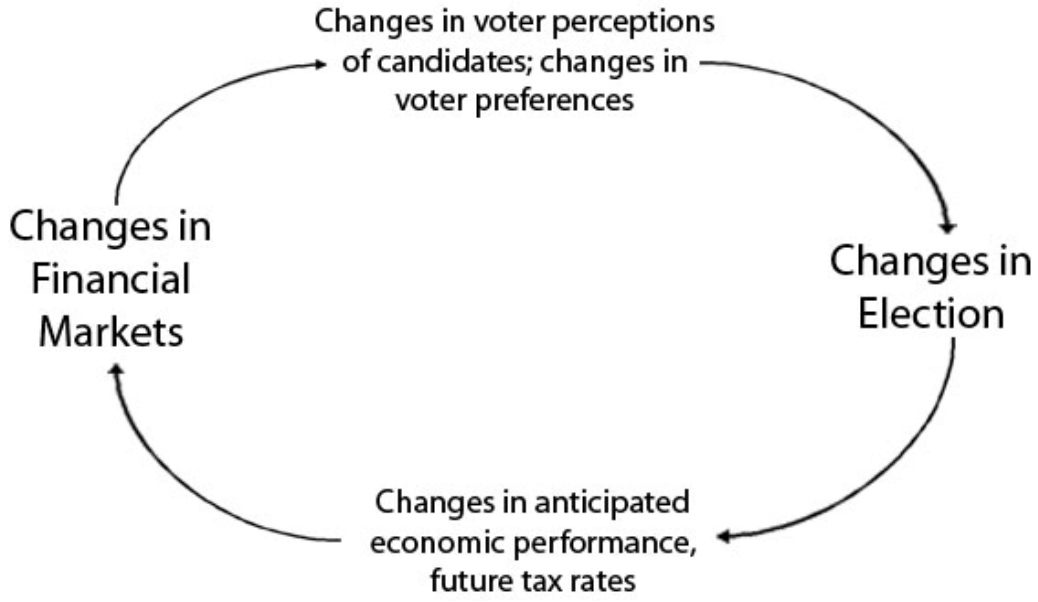

Before we can measure the impact of contextual effects on the economy, we must first isolate these dual processes. By analyzing the incorporation of election expectations into financial markets on Election Day - once voter decisions have already been made – we are able to identify the impact of the election on financial markets. After accounting for this effect, we measure the impact of various contextual effects on the election.

Empirical Methodology

In order to isolate the direction of causality for the association between financial markets and elections, we observe changes in financial market prices after voting decisions are fixed. Following the methodology proposed by Snowberg et al. (2007), we observe changes that take place during high-volume trading on Election Day as results are filtered by electronic prediction markets. By this point, all votes have been decided, but markets still have incomplete information about the ultimate election outcome. As information about the election results is revealed through exit polls and precinct reports, both prediction markets and financial markets simultaneously incorporate the information over the course of the day. Because we assume voting decisions are largely already determined, changes in financial markets are due to changes in perception about the election, rather than vice versa.

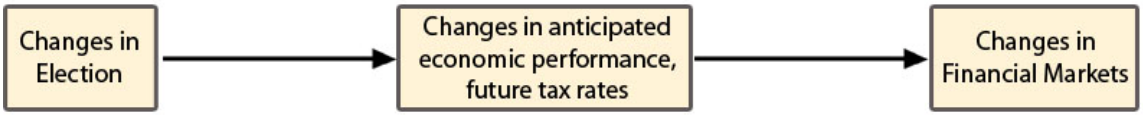

To measure this effect, we segment electronic prediction market prices and financial data over the course of the Election Day. We pool each trade that takes place into five-minute time segments. Thus, for each segment, we observe the change in price for each financial market and prediction market. Following a simple model, Equation 3.1 shows how we measure the impact of a change in prediction market prices on financial market prices:

$$\Delta FinancialMarketPrice_t = \alpha_0 + \alpha_1 * \Delta P(Victory_{jt}) + e_t \quad \textbf{(Equation 3.1)}$$

$\Delta FinancialMarketPrice$: change in financial market price at time t
$\Delta P(Victory_{jt})$: change in the probability of victory of candidate j between time period t-1 and t
e_t: error term

Data

The dataset for this portion is provided by Intrade (a leading electronic prediction market) and financial futures markets. On Intrade, we use the Obama Winner-Take-All Contract, which pays off $0 if Obama loses the election and $10 if he wins. We assume that the market price divided by the payoff price is the real-time consensus market probability of Obama winning the election. Obama was the expected winner for most of the election, and on Election Day his probability of victory increased from 91% to 99%.

Figure 3.1: Stock and Prediction Market Changes on Election Day

Figure 3.1 tracks the S&P 500 futures price and Obama Winner-Take-All prediction market probability over the course of the Election Day. Except for a late afternoon divergence, the markets seemingly track each other over the course of the day.

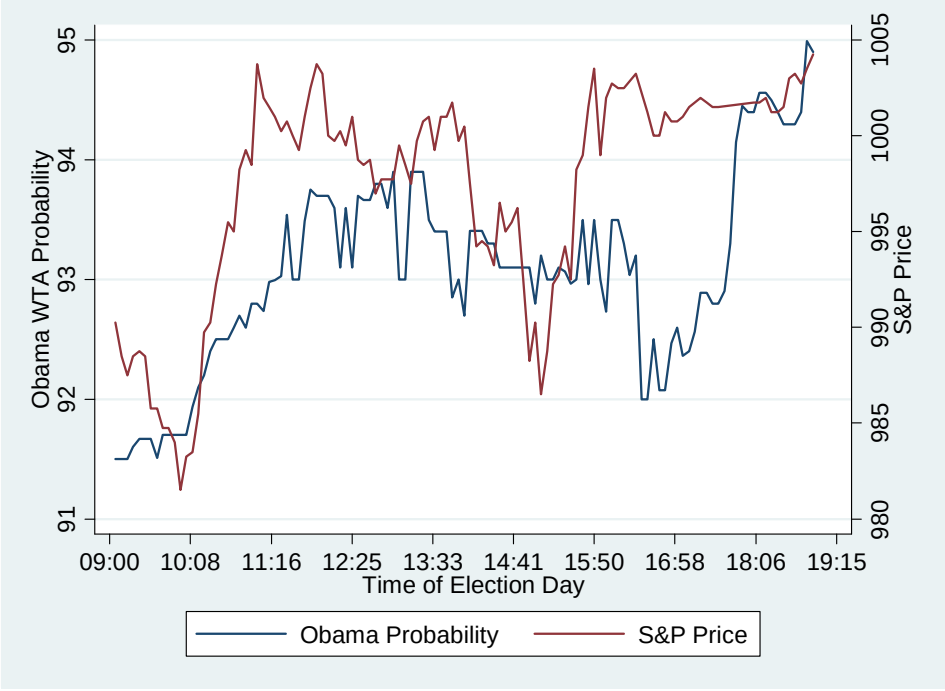

Figure 3.2: Election Changes vs. Futures Changes on Election Day

Figure 3.2 shows changes in electronic prediction market prices versus percent changes in S&P 500 futures over the course of Election Day. There is a generally positive association between the two (with significant heteroskedasticity around 0).

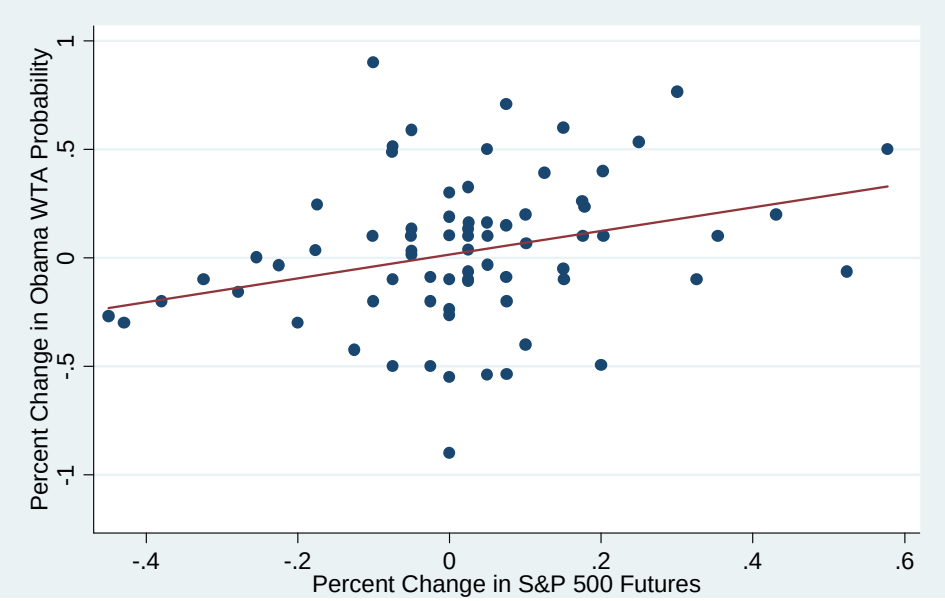

In addition to the Intrade data, we use futures market data on November 4, 2008. We track the following indices as representative financial futures: ESZ8 (S&P 500 with December 2008 settlement), CLZ8 (crude oil with December 2008 settlement), ZFZ8 (5 Year T-Note with December 2008 settlement), and ZBZ8 (30 Year T-Note with December 2008 settlement).

For each time period, we use the last trade to represent the closing price for that segment. We begin our analysis at 9:00 AM on Election Day (November 4, 2008) – when markets begin to open and information about election turnout first becomes apparent around the country. Prediction markets decline in accuracy as they approach extremes, so we end our analysis at 7:15 PM, when the implied probability first exceeds

95%. Figure 3.1 shows the movement of both the S&P 500 prices and the prediction market probability over the course of our time sample.

Results and Discussion

Applying Equation 3.1 to each of these financial indices, we find that the probability of Obama winning the election is highly associated with the S&P 500 Index, but not with any other factors. Figure 3.2 shows the relationship between lagged S&P 500 price changes and changes in Obama's election probability. Because the financial markets have greater liquidity and volume than prediction markets, we unsurprisingly

Table 3.1: Impact of Election on Financial Markets

Table 3.1 shows the results of Equation 3.1, demonstrating the effect of changes in prediction market prices on the log of the S&P 500 index on Election Day in 2008. The first column uses 5-minute time segments throughout the day (as well as a 5-minute lag), while the second column uses 30-minute segments instead. Each shows significant impact of election changes on the S&P 500. Heteroskedasticity-adjusted standard errors are in brackets. Other variables we tested had insignificant results, and they are thus omitted from these results.

Variable	(I)	(II)
ΔObama_WTA$_{t+1}$	0.001605** [0.000644]	
ΔObama_WTA$_t$		0.003270** [0.001396]
_const	0.000192 [0.000221]	.000540 [0.001037]
Observations	70	17
R-squared	.0813	0.1133
*significant at 10%	**significant at 5%	***significant at 1%

find that it incorporates information slightly more quickly than prediction markets. After testing for several possible lag lengths, we find that a 1-period (5-minute) lag is most appropriate.

Column 1 of Table 3.1 shows the results of the regression in Equation 3.1 for the S&P 500 futures index. We find that a 1% increase in Obama's probability of victory is associated with a .16% increase in the S&P 500 futures. For robustness, we now partition the Election Day into 30-minute time segments. In this case, no lag is needed, since financial markets and prediction markets both incorporated information together over this length of time. Column 2 of Table 3.1 shows that the results using these partitions confirm the strong association between election changes and equity changes. Beyond the S&P 500 index, changes in the election had no significant effect on any of our other financial indicators, and we thus omit the other results.

By confining the sample to Election Day, our methodology allows us to interpret the association as a *causal* relationship: markets rose because Obama's likelihood of victory increased. There are three primary reasons why this may be the case. First, voters may anticipate real economic changes due to Obama's policies. A large stimulus package and an intelligent economic team may increase expected future earnings for companies in general and thus boost equity prices. Secondly, rather than increasing economy-wide growth in general, Obama's policies may be helpful to specific companies in the S&P 500 Index. Through company- or sector-specific bailouts, Obama's policies could be beneficial to companies that are listed on the S&P Index and support their stock prices. Finally, changes to tax policies could make equities relatively more attractive on an after-tax basis, increasing prices now as voters predict future changes (given Obama's

stances on capital gains taxes, this does not seem like the most compelling explanation of the three). A combination of these factors made Obama's election a positive factor for equity prices.

In comparison with past years, Obama's election appears to have greater impact on financial markets than most elections, and he is one of the rare Democrats whose victory was causally associated with a stock market increase on Election Day. Using a similar methodology, Snowberg et. al. (2007) found that a 1% increase in Bush's probability of victory in both the 2000 and 2004 election was associated with a .015% increase in the S&P 500 index, and that favorable election results for Republicans generally boosted equity markets. The relative size of the shock may be due to the fact that the economy was likely more central to this election than others, and market participants expected the winner to affect economic policy more substantially than in a typical election. Indeed, the financial crisis in 2008 drew frequent media comparison to the Great Depression, and the extent that a 1% change in Obama's probability of winning affected financial markets rivals the impact that occurred in the 1932 election. According to Snowberg et. al. (2007), betting markets indicated that a 1% increase for Franklin Roosevelt in 1932 was associated with a .25% decrease in the stock market. The large size of the effect in 2008 indicates that – like in 1932 – investors likely expect significant economic policy changes to be made under the new president.

Part II: Financial Effects on the Election

We now apply the results from Equation 3.1 to identify the inverse effect: the impact of financial markets on the 2008 election.

Empirical Methodology

In order to measure the effects of financial markets on elections, we expand on Equation 2.3, which proposes that the election is shaped by a combination of electoral context and candidate particulars (charisma, views, likeability, etc.). Thus, the probability at time *t* that candidate *j* will win the election is:

$$P(victory_{j,t})=\beta_0 * Candidate\_Particulars_{j,t} + \beta_{1,j} *Context_t + u_{jt} \quad \textbf{(Equation 3.2)}$$

$P(Victory_{jt})$: probability of victory of candidate j at time t
$Candidate\_Particulars_{j,t}$: impact of candidate-specific traits, decisions, and policies on the election for candidate j at time t
$Context_t$: relevant contextual factors, such as S&P 500 prices, economic indicators, oil prices, etc. at time t
u_{jt}: error term

$Context_t$ is a vector of contextual factors (such as equity, oil, and bond prices) that describe the state of the world at time *t*, and $\beta_{1,j}$ is a vector of their weights in the election for candidate *j*. We assume that candidate particulars are orthogonal to the context of an election, and they can thus be treated as unobserved fixed effects. This implies that we can measure changes in electoral probabilities as:

$$P(victory_{j,t+1})- P(victory_{j,t})= \gamma_0 + \gamma_{1,j} (Context_{t+1} - Context_t) +\varepsilon_{jt} \quad \textbf{(Equation 3.3)}$$

Or: $\Delta P(Victory_j) = \gamma_0 + \gamma_{1,j} *\Delta Context_j + \varepsilon_j$

$P(Victory_{jt})$: probability of victory of candidate j at time t
$Candidate\_Particulars_{j,t}$: impact of candidate-specific traits, decisions, and policies on the election for candidate j at time t
$Context_t$: relevant contextual factors, such as S&P 500 prices, economic indicators, oil prices, etc. at time t
ε_{jt}: error term

Thus we model the effect of a change in context on the change in probability of victory from one period to the next. In order to isolate the directionality of this effect, we adjust

the values for $\Delta Context_t$ by the predicted change in economic variables. Using Equation 3.1, we construct a variable called Predicted_ $\Delta Context_{jt}$, defined below as the unexpected change in context, and plug it into our regression equation.

(Equation 3.4)

$$\Delta P(Victory_{jt}) = \lambda_{0,j} + \lambda_{1,j} * (\Delta Context_t - Predicted\_\Delta Context_t) + \varepsilon_{jt}$$

Where: $Predicted\_\Delta Context_t = \alpha_0 + \alpha_1 * \Delta P(Victory_{jt})$

$P(Victory_{jt})$: probability of victory of candidate j at time t
$Context_t$: relevant contextual factors, such as S&P 500 prices, economic indicators, oil prices, etc. at time t
$Predicted\_\Delta Context_t$: expected changes of financial markets due to election, as calculated above
ε_{jt}: Error term

In this equation, we make two strong assumptions. First, we assume that the impact of the election on equity prices is linear; if V is the net present value of a candidate winning to the S&P 500 and p represents his probability of winning, p*V should be the premium currently associated with the candidate for a risk-neutral market. Secondly, we assume that the perception of candidate's value is constant over time; in reality, V may fluctuate as markets (like individuals) respond to new conditions and learn more about candidate priorities. Our results will allow us to relax each of these assumptions.

Since equity prices are the only financial market in which we found a significant effect from the election, this is the only contextual factor that must have a predicted value associated with it. By accounting for the reverse causality, we are now able to directly identify the impact of financial markets on the election.

Figure 3.3: Probability of Obama's Victory over Course of Election

Figure 3.3 shows the electronic prediction market consensus probability for Obama's victory over the course of the election. As can be seen, his likelihood of success fluctuates around 60% until a steep increase in mid-September.

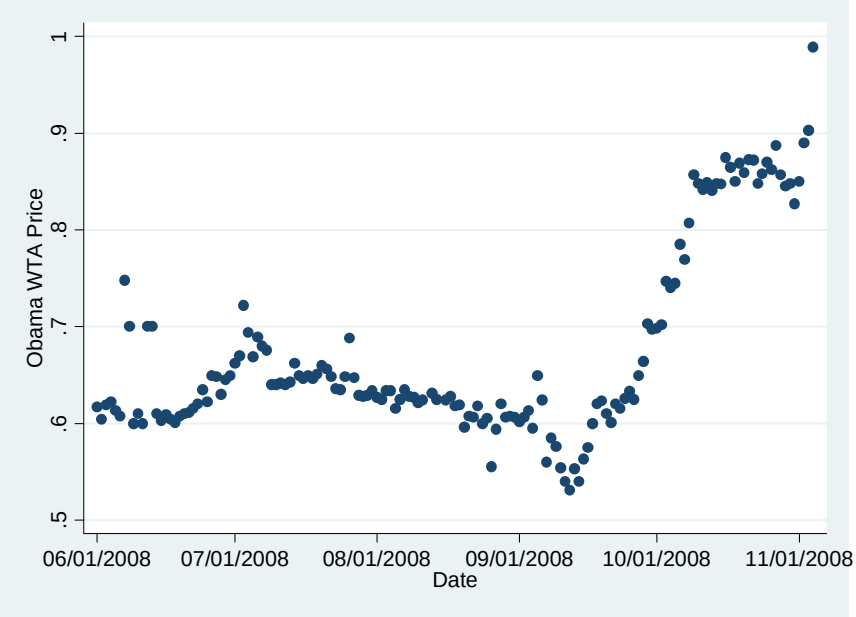

Data

Part II uses data from the Iowa Electronic Markets and financial future markets. While Intrade has larger volume, there is evidence that a rogue trader attempted to manipulate results for a portion of the election, so results from Iowa Electronic Markets is more consistent for the longer time horizon.[1] The two markets are otherwise nearly identical, due to cross-market arbitrage opportunities. As before, we use the Obama

[1] According to David Rothschild and Justin Wolfers: "Over [several] weeks we've observed a pattern of large orders for Sen. McCain on Intrade…executed at times when liquidity is particular scarce. These orders have caused markets to shift sharply, often against the broader political narrative… The suspicious behavior…involves large purchases of Sen. McCain's stock executed at times when most traders wouldn't be active. These price changes are unusually large and occurred during periods of relative political calm. By comparison, neither of the vice-presidential announcements caused comparable market shifts on Intrade. And prediction markets assessing the fate of important swing states did not mirror these movements in Sen. McCain's overall share price. Alternative prediction markets, including BetFair, British sports books or play-money markets, also failed to register similar movement for the Arizona senator."

Winner-Take-All Contract, assuming that the market price divided by the payoff price is the real-time consensus market probability of Obama winning the election. We analyze the period from June 1, 2008 to November 4, 2008. Figure 3.3 shows the progression of predicted probabilities over time.

We track the following financial indicators: the S&P 500 Index, the AMEX Oil Index, the Chicago Board of Exchange Volatility Index, the 5-Year T-Bill price, and the 30-Year T-Bill price. Additionally, in order to measure real economic changes rather than just financial changes, we track the Intrade Recession Index. The Intrade Recession Index is a winner-take-all market that tracks the market consensus probability of recession at any point in 2008. Figures 3.4-3.9 show the values of each of these indicators.

For each of these datasets, the price for each day is defined as the closing price, or the last trade that takes place. To capture longer-term effects, each difference we use is the rolling 7-day price change.

Results and Discussion

Our results indicate that economic and financial changes played a substantial role in shaping the 2008 election. Figure 3.10 shows the relationship between the election and the S&P 500, and Figure 3.11 shows the relationship between oil prices and the election. In both cases, there appears to be a strong negative association between the two (not accounting for other factors).

Figure 3.4: S&P 500 Price during Election

Figure 3.4 shows the change of the S&P 500 index value over the course of the election, indexed at 1 on June 1, 2008. The price declined rapidly beginning in mid-October, and the index lost approximately 30% of its value over the course of the election.

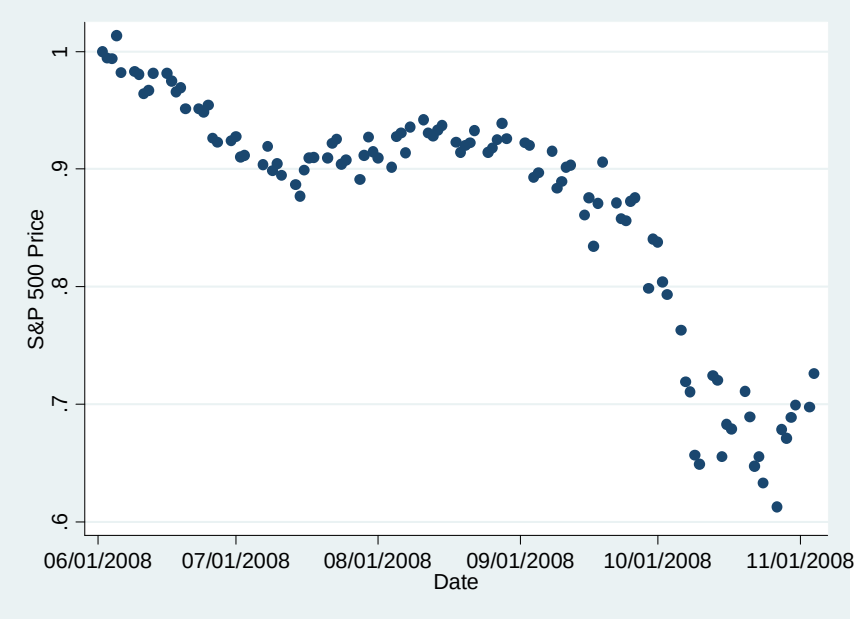

Figure 3.5: Oil Prices during Election

Figure 3.5 shows the change of the Amex Oil Index value over the course of the election, indexed at 1 for the price on June 1, 2008. The price declined quickly throughout the election, losing approximately 40% of its value.

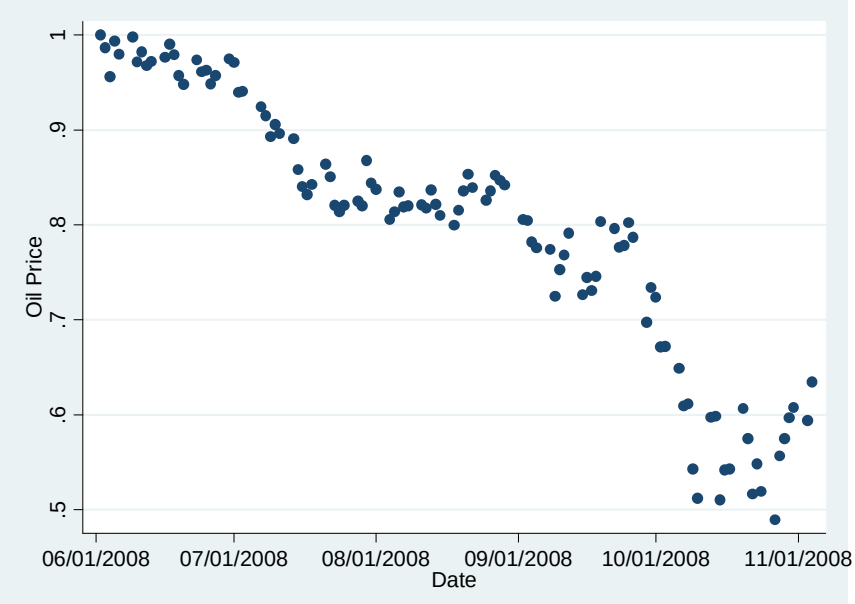

Figure 3.6: Volatility during Election

Figure 3.6 shows the change in the VIX volatility index during the election, indexed at 1 for the price on June 1, 2008. The price spiked beginning in mid-October, and volatility was nearly four times its June 1 value at its peak.

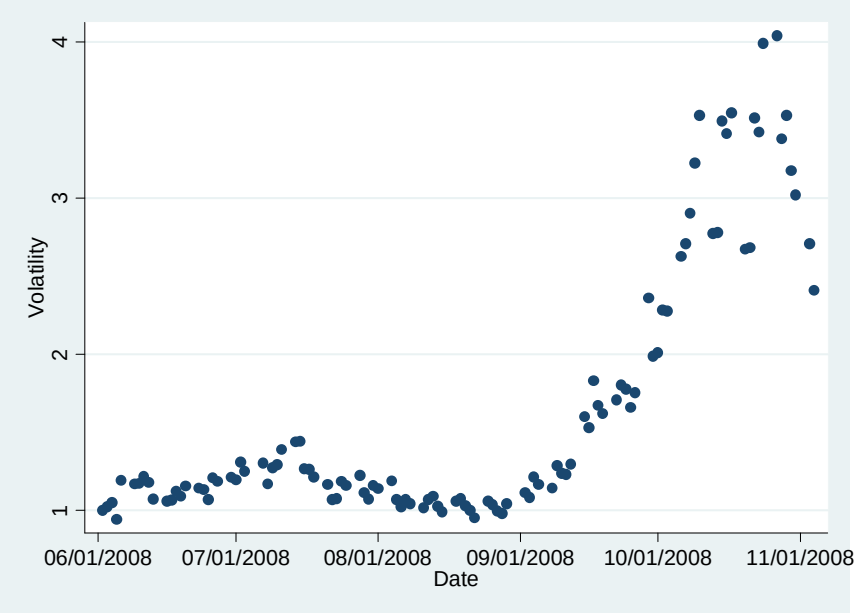

Figure 3.7: 5-Year Treasury Price during Election

Figure 3.7 shows the change of 5-year treasury prices over the course of the election, indexed by the price on June 1, 2008. The price shows a gradual decrease over the course of the election.

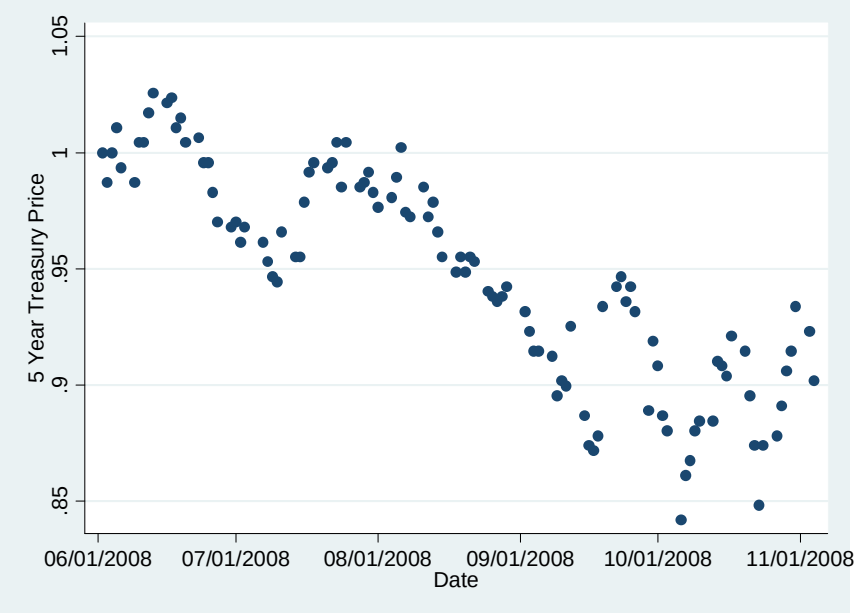

Figure 3.8: 30-Year Treasury Price during Election

Figure 3.8 shows the change of 30-year treasury prices over the course of the election, indexed by the price on June 1, 2008. The price shows a gradual decrease over the course of the election.

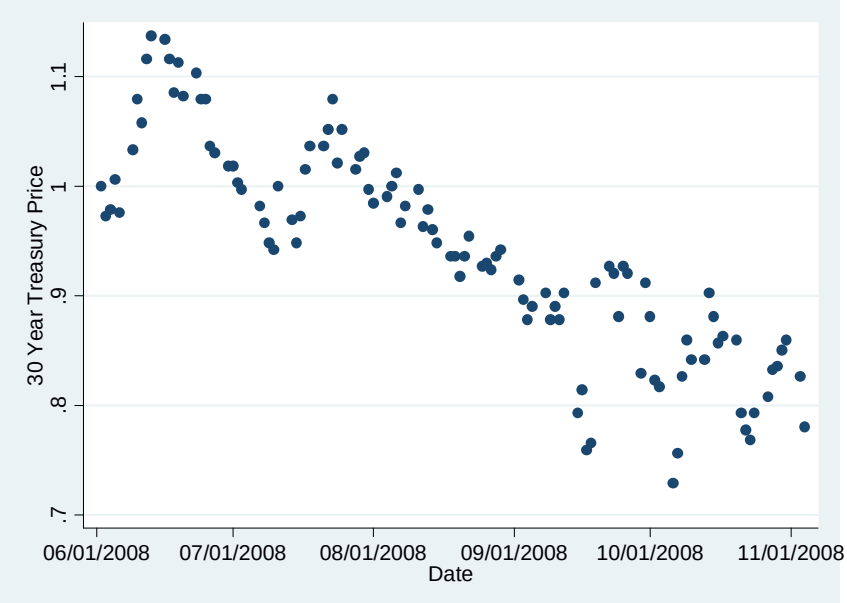

Figure 3.9: Recession Probability over the Election

Figure 3.9 shows the change of the Intrade Recession Index over the course of the election. The index is traded on Intrade, and it pays off $10 if a recession occurs in 2008 and $0 otherwise. Prices are indexed at 1 for the price on June 1, 2008. The price shows a sharp increase beginning in mid-September.

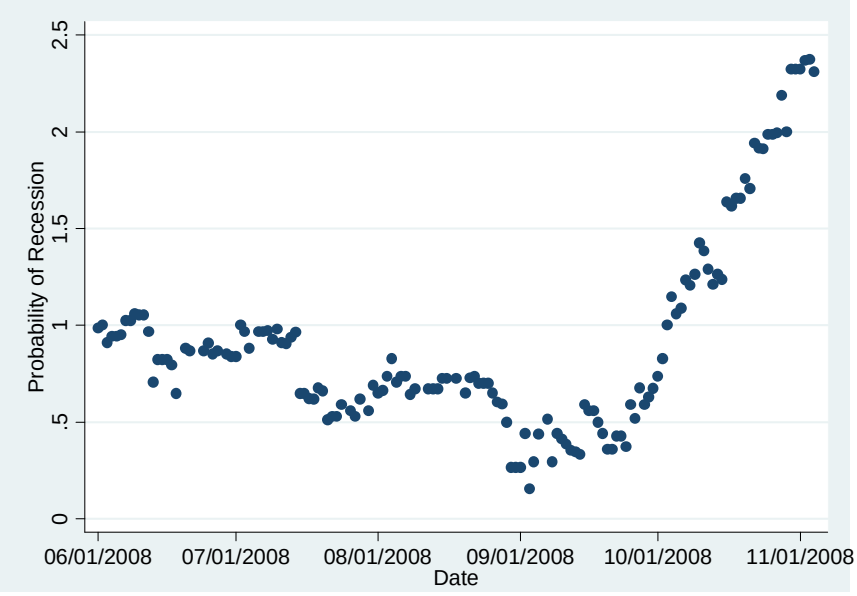

Figure 3.10: Association between S&P 500 and the Election

Figure 3.10 compares the rolling 7-day change of Obama's probability of victory with changes in the S&P 500. There is a negative association between these two variables (with heteroskedasticity around 0).

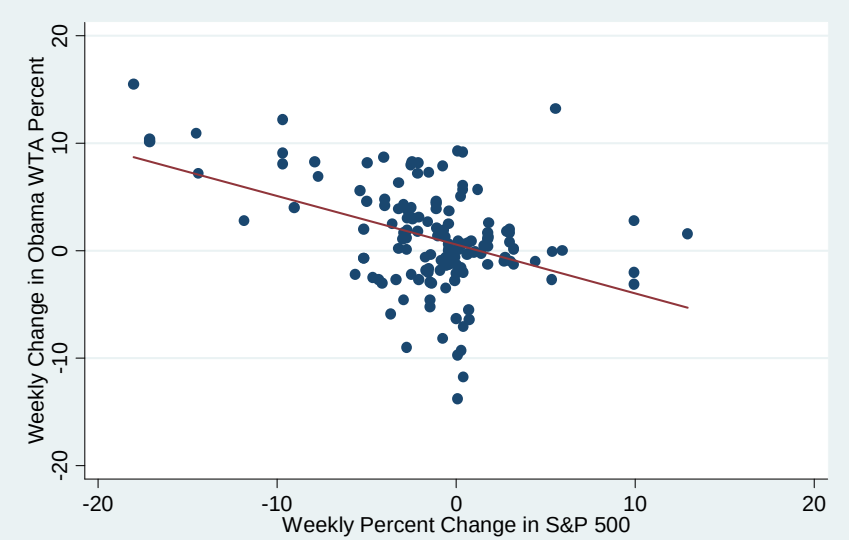

Figure 3.11: Association between Oil Prices and the Election

Figure 3.11 compares the daily change of Obama's probability of victory with changes in the Amex Oil Index. There is a negative association between these two variables when we do not account for any other factors (such as equity changes).

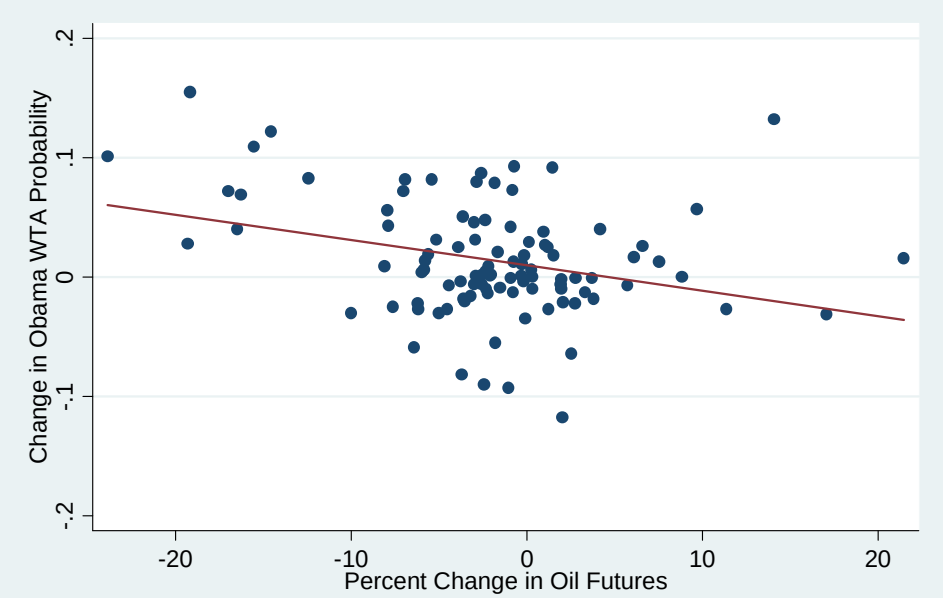

Once we account for several financial factors simultaneously, our results provide a more complete picture of each variable's impact on the election. Column 1 of Table 3.2 shows the results of Equation 3.4 using the estimate of the impact of political changes on the election from Part I. We find that a 1% increase in equity markets is associated with a 0.929% decrease in Obama's probability of victory, a 1% increase in oil prices is associated with a 0.27% increase in Obama's probability of victory, and a 1% increase in

Table 3.2: Impact of Financial Markets on the Election

Table 3.2 shows the results of Equation 3.4, demonstrating the effect of changes in financial markets on Obama's probability of victory over the course of the election. Log transformations were applied to each variable. Log_Equity_Deviation is defined as the change in the log S&P 500 value from its expected value based on the impact of the election on it. Each column uses a different assumption for impact of the election on the S&P 500. Column 1 uses the expectation from Table 3.1, Column 1, assuming a 1% increase in Obama's probability causes a .16% increase in the S&P 500. Columns 2, 3, and 4 assume 0%, 0.1%, and 0.2% impacts respectively. Each result shows a significant impact of financial factors on the election. Heteroskedasticity-adjusted standard errors are in brackets. Insignificant variables are omitted.

Variable	(I)	(II)	(III)	(IV)
Log_Equity_Deviation	-0.929***	-0. 658***	-0.858***	-0.959***
	[0.121]	[0.129]	[.123]	[.120]
ΔLog(Oil)	0.270**	0.094	0.218*	0.295***
	[.110]	[0.108]	[0.111]	[0.108]
Δlog(5Year_Bond)	0.311***	0.282**	0.307***	0.310***
	[0.104]	[0.113]	[0.107]	[0.103]
_const	0.005	0.007	0.006	0.004
	[0.004]	[0.005]	[0.004]	[0.004]
Observations	99	99	99	99
R-squared	0.399	0.219	0.330	0.442

*significant at 10% **significant at 5% ***significant at 1%

bond prices is associated with a .31% increase in Obama's probability of victory. Once we account for each of these factors, all other variables we tracked are insignificant.

The very strong negative association between equity prices and Obama's probability of victory confirm our expectation that the financial crisis had a significant impact on the election outcome. From September 5 to October 15, the stock market plunged 27%, while Obama's probability of victory surged from a near tie to a dominant lead. The concurrence of these two events was no mere coincidence; both the *saliency* and *clientele hypotheses* indicate why this might be the case. Under the *saliency hypothesis*, voters may have blamed the incumbent Republican Party for weak markets and shifted ideological views away from a *laissez faire* approach. The new popularity of regulation would increase the typical voters' expected agreement with Obama's policy views ($v_{i,Obama}$), boosting his odds of success. Additionally, under the *clientele hypothesis*, the economic crisis likely increased the weight the average voter assigned to economic policies in the election (i_i). Voters shifted their voting priorities away from social and foreign policies towards economic concerns, favoring the Democrats in this case. The combined realignment of voter priorities and general perception of regulation strongly boosted Obama's campaign during a period of great economic uncertainty.

The positive association between bond prices and Obama's likelihood of victory is less straightforward to interpret, accounting for other financial factors. An increase in bond prices may indicate a decrease in real interest rates, inflation expectations, or government debt default risk, and there is no clear indication which of these may be

favorable for Obama.[2] Perhaps the most compelling explanation is that fears of a deflationary spiral boosted Obama's candidacy while decreasing expected inflation rates. Alternatively, under the *clientele hypothesis*, if Republicans are more trusted to fight inflation, decreased inflation fears caused a decrease in the weighted importance of this as a factor to vote on (i_i) and thus boosted Obama.

Surprisingly, while changes in the Intrade Recession Index are strongly correlated with Obama's probability of victory, we find that changes in the index have no significant effect on the election once we account for financial indicators. This implies that financial changes were more significant than real economic changes – a contrast to the typical Fair model use of real economic factors. However, this is likely due to the unique nature of this year's election at the center of a financial crisis in which all eyes were turned towards the stock market.

The positive association between oil prices and Obama's probability of victory explains the impact that a key source of campaign rhetoric had on the election. During the summer's period of high oil prices, each candidate criticized the other's policies as ineffective in alleviating the dependency on foreign oil, and gas prices became a focal point of the campaign. According to this result, the rise in prices helped Obama, though the subsequent decline hurt his campaign. High gas prices were the initial catalyst in making the economy a central point in the election, and they boosted Obama's candidacy from the beginning of the campaign.

[2] One method of isolating inflation expectations and interest rates would be examining inflation-indexed TIPS bonds. However, over this time period, TIPS prediction may be inaccurate, as TIPS does not adjust for *deflation*, which became a reasonable risk during the financial crisis.

Concluding Remarks

Taken in aggregate, the large economic shift that occurred was a decisive factor in shaping the election. For the period from September 5 to October 15, the rapid change in each of these areas had major effects on the election and dramatically shaped its ultimate results. The major increase in Obama's probability of victory we predict from the decline in equity markets was partially offset by the decreases in his probability of victory due to increased bond and oil prices, but the aggregate effect was a predicted 15.9% increase in his probability of victory, explaining much of the actual 19.8% shift that occurred over that time frame.

Beyond the analysis performed in this chapter, the methodology represents a broader implication of the rapid growth of electronic prediction markets: the ability to quantitatively dissect causal effects within a single election – even in an election in which one candidate had a comfortable lead throughout. Without a continuous source of information on election status, the bi-directional nature of the interaction between the economy and elections could not be disentangled, and changes in priorities that occur during an election year would go unnoted. Prediction markets enable the possibility of an event study on an election, and they empower political economists and political scientists to better understand the relationship between elections and economic changes.

Chapter IV

Is Context Everything?

Measuring the Impact of Campaign Events and Advertising

In this chapter, we turn to the impact of campaigning on the 2008 U.S. presidential election. During electoral campaigns, candidates boost their appeal to voters by distributing information in the form of campaign events and advertising. Using electronic prediction market prices to control for the probability of victory at a particular instance of time, we observe how the candidate's choice of resource allocation in that time period affects the election results across states. We find that a 1% increase per electoral vote in the share of campaign events results in a 1.26% increase in Obama's final margin of victory in a particular state. Similarly, a 1% increase in McCain's share of events per electoral vote results in a 0.38% decrease in Obama's final net margin of victory in a particular state. We find that campaign advertisements have no significant local-level effect. We compare these results with the impact of other determinants of election outcome in light of our general model of elections.

Introduction

In the previous chapter, we focused on the impact of economic context on the election. To a candidate, economic and financial factors are exogenous, and the candidate is left with no control over the election's outcome. In this chapter, we examine election factors that a candidate can control, and the actions a candidate might take in trying to win an election.

An extreme interpretation of the Fair model is that of economic determinism, in which an election is purely controlled by its environment and campaigns have no impact. Fair's model, which can predict the outcome of an election extremely effectively with just knowledge of economic context and the incumbency, does indeed indicate – and our previous chapter confirmed – that context plays a crucial role in an election. As Shaw (1999) notes, "although presidential campaigns have been mythologized in literature and cinema, most theories of elections relegate them to a secondary role, presuming they have little effect on outcomes." Within a slight margin of error, an election can be fully predicted based on contextual events – and campaigns are seen as mere byproducts of an already-determined election.

Yet this analysis contradicts an apparent reality: campaigns are major events in which significant resources are invested. If they have no impact, why do rational individuals invest in them? In this chapter, we consider two ways in which candidates can affect the election: through hosting events and buying advertisements. Measuring the impact of each of these activities is, however, quite difficult. There are two main sources of empirical problems. The first challenge, as Shaw (1999) states, is a lack of "data on the allocation of resources and voters' preferences." A second, more subtle, problem is

due to endogeneity biases. As Levitt (1994) notes, most studies do not "adequately control for inherent vote-getting ability across candidates." Candidates who are more likely to win elections may receive more money due to their increased popularity or due to donors' attempts to buy future goodwill, so measurements of the impact of campaign-wide spending are typically biased. State-level cross-sectional analyses also suffer from bias, as candidates choose their resource allocation strategy based their initial probability of winning.

With electronic prediction markets, we are able to control for this bias. By controlling for the starting status of the election across states, we measure the impact of campaign allocation of resources on an election. We find that while campaign events have a significant impact on the election outcome, campaign expenditures do not have any significant local-level effects. Among events, Barack Obama's have much greater impact than those hosted by John McCain.

The assumptions made in this chapter are much stronger than the assumptions made in Chapter III, and the results – while significant – are more limited in scope. Nonetheless, this chapter illustrates how electronic prediction markets are useful in understanding campaign resource allocation, and we discuss how further use of prediction markets could in fact enable us to test the assumptions we make.

This chapter is structured as follows: first, we model candidate impact on voters in light of our general model of voter behavior; second, we consider an empirical methodology using electronic prediction markets to determine the impact of candidate choices; finally, we discuss the implications of our results.

Model and Methodology

While most of the factors impacting our general model of elections, Equation 2.3, are exogenous to candidates (such as voter taste, context, and partisan signal effects), there are several factors that the candidates do control. In particular, the candidates can – to a limited degree – control their personal characteristics (p_j) as well as the correlation between their beliefs and the median voter's (r_{ij}). There are two ways in which candidates can shape these factors: through direct changes in beliefs or through informational changes.

We will now model this relationship more formally as the vector of revealed belief correlation (r_{ij}) from Equation 2.3. Each element of the vector represents the extent to which a candidate's past views indicate agreement with a voter on a particular topic. There are two components that determine the level of revealed belief correlation between a candidate and a voter: the perceived extent of agreement (ρ_{ij}) and the degree of certainty about the extent of agreement (d_{ij}). If a voter is less certain about a candidate's opinion on a particular topic, d_{ij} will be lower and the voter will thus discount her perceived degree of agreement due to the uncertainty. We model these dynamics in Equation 4.1:

$r_{ij} = r(\rho_{ij}, d_{ij})$ **(Equation 4.1)**
Where: $\partial r_{ij}/\partial \rho_{ij} \geq 0$
And: $\partial^2 r_{ij}/\partial d_{ij}\partial \rho_{ij} \geq 0$

r_{ij}: perceived impact of revealed opinions on future agreement
ρ_{ij}: perceived correlation of past opinions between voter i and candidate j
d_{ij}: degree of certainty about extent of agreement between voter i and candidate j

We assume that a candidate has limited control over the degree of correlation with the median voter (ρ_{ij}), which can be altered by altering policy views. The candidate has more

control over the degree of certainty the voter has about their level of agreement (d_{ij}), which can be altered by changing the voters' awareness about particular policy positions.

As a more concrete example, consider the issue of taxation. While there is general disagreement among voters about whether taxes should be raised for wealthy individuals, a much more popular appeal to a given voter is that taxes should be cut within her own tax bracket. The model indicates that a candidate who plans a tax increase for the top 1% with a simultaneous tax decrease for the bottom 99% has two ways of gaining popularity: changing opinions to boost agreement with the median voter (ρ_{ij}), or focusing on increasing the informational certainty the average voter has about the impact of the tax change on the broader 99% by drawing attention to it through speeches and commercials. A voter, with time and informational constraints, does not know all components of a candidate's platform, and thus discounts issues in which the candidate's stances are less known. The candidate can circumvent this by increasing the median voter's certainty on points of agreement (or, alternatively, increasing the voter's certainty about points of disagreement with the opponent).

However, candidates incur a cost to distributing information that focuses voters on key points of agreement. We assume that the probability of a particular candidate winning a state is a function of the candidate's resources allocated towards distributing information in that state, the opponent's resources allocated, the initial probability of winning the state without allocating any additional resources (which varies from state to state based on the v_{ij} and i_i of the median voter in that state), and national effects. Thus we model the probability in period t+1 of Obama winning a given state as:

$P_{j,s,t+1} = G(o_{jst}, o_{kst}, P_{j,s,t}, N_{j,t})$ **Equation 4.2**

$P_{j,s,t+1}$: probability at time t+1 that candidate j will win in state s
o_{jst}: resource/money allocated by in state s by candidate j at time t
o_{kst}: resource/money spent in state s by opponent candidate k at time t
$P_{j,s,t}$: probability at time t that candidate j will win in state s
$N_{j,t}$: national fixed effects on candidate j at time t

A critical assumption of this model is that the additional resources are all allocated immediately. If money is continuously distributed, it creates an endogeneity bias, as candidates deploy money based on the probability in period t+1 as well. Candidates, seeking to maximize their probability of winning the election, look for states where the impact of their additional resources can have maximal impact and thus base their decisions on probabilities. While the function may vary from candidate to candidate

Figure 4.1: Strategic Allocation of Campaign Spending for John McCain

This figure compares the advertising spending allocation for John McCain in swing states over the week of September 28, 2008-October 4, 2008 with electronic market prediction price in each state on September 27, 2008. In swing states where there was a higher probability of Obama winning, McCain was more aggressive with his campaign expenditures. This confirms the endogeneity bias associated with many estimates of campaign spending impact.

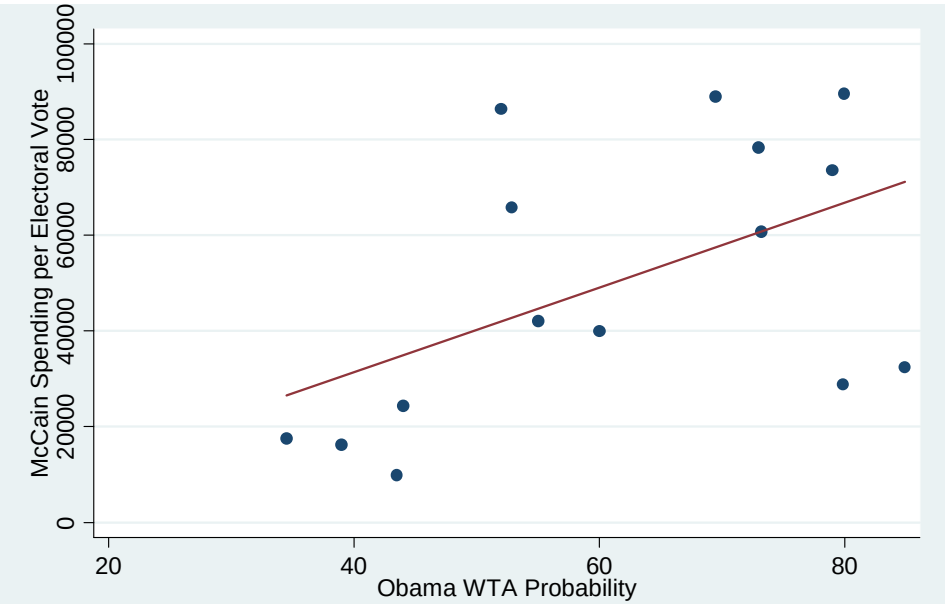

based on game theoretical considerations, each candidate makes their allocation decisions as a function of their probability of victory. Figures 4.1 and 4.2 show the extent to which probability affects candidates' allocations, even within relatively close states; McCain spends more in swing states where he is further behind, while there is no clear pattern to Obama's spending patterns.

This endogeneity bias due to allocation decisions being made based on probability of victory creates an intriguing dilemma. Since spending allocations are made on a continuous time horizon, observing the allocation of campaign resources over any finite time period will generate a bias, as candidates may spend more in states where they are performing more or less successfully already. If the candidates select their spending

Figure 4.2: Strategic Allocation of Campaign Spending for Barack Obama

This figure compares the advertising spending allocation for Barack Obama in swing states over the week of September 28, 2008-October 4, 2008 with his electronic market prediction price in each state on September 27, 2008. There is no clear association in Obama's allocation pattern.

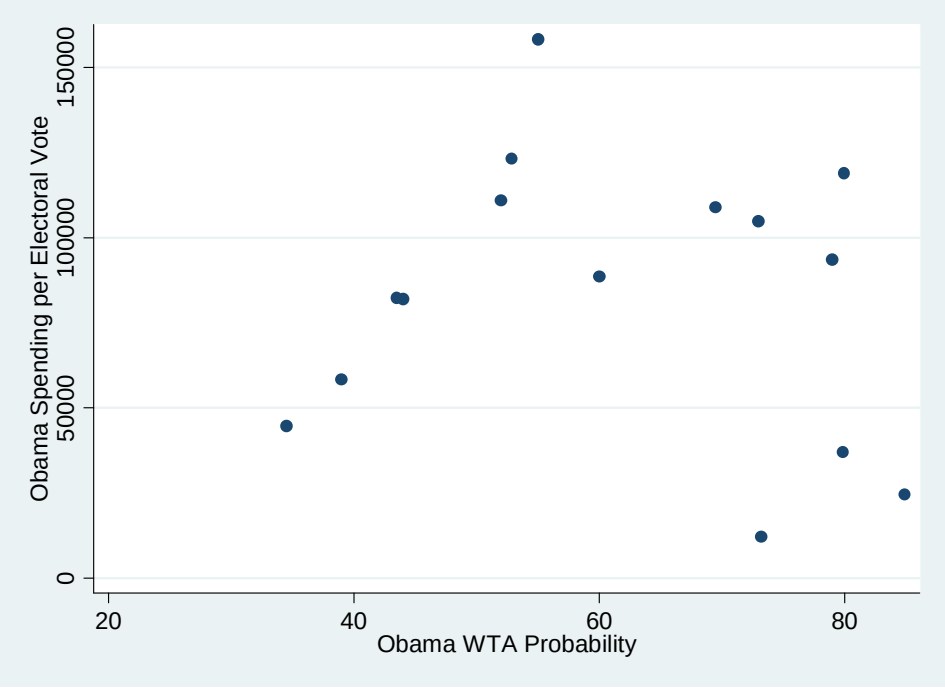

allocations towards the end of the time period, spending will reflect the probabilities towards the end, biasing our result. We circumvent this by measuring changes in the election over a long time horizon but resource allocation over a short period. We use the first week's spending allocation as a proxy for spending throughout the entire period. If we assume that very minimal changes would occur to spending allocations after the first week of our period and that allocations of resources over this week are unbiased estimators of the future share of resources allocated, then we can replace Equation 4.2 with a revised version.

$P_{j,s,f} = G(o_{jst}, o_{kst}, P_{j,s,t}, N_{j,t})$ **Equation 4.2**

$P_{j,s,f}$: probability at end of the election (period f) that candidate j will win in state s
$o_{j,s,t}$: resource/money allocated in state s by candidate j at time t
$o_{k,s,t}$: resource/money spent in state s by opponent candidate k at time t
$P_{j,s,t}$: probability at time t that candidate j will win in state s
$N_{j,t}$: national fixed effects for candidate j at time t

This creates a new dilemma. Because the 2008 election was not particularly close, by the end of the election most previously-swing states were trading at very high probabilities of Obama victory. This leaves little room for calculation. Instead, we use the actual vote spread in the election, which provides a greater range of values.

We now attempt to formulate our linear model. There are two main types of resources that can be expended in order to provide positive information to voters. The first is advertising, which can be measured by expenditures. The second method is through campaign events, which both directly inform voters and generate additional media attention in that state.

Equation 4.3

$$Vote\_Spread_s = \beta_0 + \beta_1 * ObamaEventShare_{st} + \beta_2 * McCainEventShare_{st} + \beta_3 * ObamaAdvertisingShare_{st} + \beta_4 * McCainAdvertisingShare_{st} + \beta 5 * P_{st} + \varepsilon_s$$

β_0: constant term that includes national effects
$Vote\_Spread_s$: percentage points by which Obama beat McCain in state s
$ObamaEventShare_{st}$: share per electoral vote of Obama's events allocated to state s over time period t
$McCainEventShare_{st}$: share per electoral vote of McCain's events allocated to state s over time period t
$ObamaAdvertisingShare_{st}$: share per electoral vote of Obama's advertising budget allocated to state s over time period t
$McCainAdvertisingShare_{st}$: share per electoral vote of McCain's advertising budget allocated state s over time period t
P_{st}: initial probability that Obama would win state s
ε_s: error term

Equation 4.3 models the final vote spread in a given state as a linear function of the share of campaign spending and advertising by each candidate over a fixed time period, accounting for the initial probability Obama would win the state.

Three major assumptions are made in this model. First, as we already identified, we assume a candidate's allocation of resources in time period t is an unbiased estimate of all resource allocation after time period t throughout the remainder of the election; spending may rise and fall in particular states, but we assume these changes are uncorrelated with the resource allocation.[1]

A more subtle assumption is the strict efficiency of our metric for initial probability ($P_{j,s,t}$). We assume that the probability is the only relevant factor (along with the opponent's response) that the candidate considers when determining the per-vote share of resources that they allocate. For our estimates to be unbiased, any other factors that are considered must have no expected correlation with the error term. Based on the core premise of this paper – that electronic prediction markets provide the best known

[1] This assumption would be testable given a complete data set. Unfortunately, no data set for 2008 overall campaign spending currently exists.

estimates at any given time for the probability of a particular outcome – our assumption is that campaign staff do not have better methods of calculating probabilities, and – if they do – that they would participate in prediction markets themselves.[2] This, in fact, is the great strength of using electronic prediction markets for this analysis; if polls were used as a metric of probability at any given time, campaigns would have better information than the model incorporates, generating a bias.

Finally, in designing this methodology, we make a strong assumption that the markets are unaware of future resource allocation decisions, and the expected impact of future resource allocation is not already incorporated into the market price. That assumption can be relaxed. Let $E(o_{jst})$ denote the expected resources used by candidate j in state s. Presumably, o_{jst} is positively correlated with $E(o_{jst})$, which is one of the constituent factors within P_{jst} in our model. If we could omit $E(o_{jst})$ from our model, the magnitude of the impact of m_{js} would increase. Thus our estimates are in fact a *lower bound* of the actual impact on a state of a candidate allocating additional resources.

Data

In this chapter, we analyze campaign spending data, campaign event data, prediction market prices, and final vote spreads from a variety of sources and time periods.

Publically available campaign spending data catalyzed our choice for both our time range and the states within our sample. The Wisconsin Ad Project – a program at the University of Wisconsin that tracks campaign spending – released spending data for

[2] Unlike financial markets, most electronic prediction markets *encourage* insider trading, as it boosts the accuracy of the prediction. In many instances, there is a pre-announcement effect as insiders trade for informational arbitrage. This was particularly notable before Larry Summers' resignation from Harvard's presidency, when markets spiked a couple hours before the story was released to the *Harvard Crimson*.

the fifteen closest "swing states" for the week of September 28, 2008 to October 4, 2008. We use aggregate spending data for advertisements purchased directly by the candidates.

To measure the number of campaign events, we use the *Washington Post's* Campaign Tracker for the period of September 28, 2008 to October 4, 2008. We tally the events that took place in each state in which the presidential candidate, the candidate's spouse, or the vice-presidential candidate participated. We omit required events, such as party functions or debates.

To calculate the initial probability in each state, we utilize state-by-state electronic prediction market prices provided by Intrade. Intrade has recently attempted to encourage trading on a state-by-state basis, in which winner-take-all contracts predict the probability of victory for a given state. During our measurement period, swing states generally had at least several trades per day, providing a reasonably accurate set of day-to-day predictions on the election outcome. We select the last trade on September 27 (the day before our advertising period began) as the representative value of the initial probability. For robustness, we also test the data with the September 21 closing price as our initial probability, and we achieve similar results.

Vote spreads for each state were taken from the *New York Times* President Map.

Results

Our results indicate that candidate events play a significant – but not overwhelming – role in the election. Table 4.1 shows the results of Equation 4.3. We find that a 1% increase in the share of events per electoral vote results in a 1.26% increase in Obama's final net margin of victory in a particular state. Similarly, a 1% increase in McCain's share of events per electoral vote results in a 0.38% decrease in Obama's net margin of victory.

Table 4.1: Regression Results

This table shows the results of Equation 4.1, demonstrating the effect of an increase in the per-electoral-vote share of events and advertising in each state during the week of September 28 on the final election spread. An increase in Obama's share of events in a given state by ELECTORAL_VOTES % is associated with a 1.26% increase in the vote spread, where ELECTORAL_VOTES is the number of electoral votes in that state. Similarly, an increase in McCain's share of events by ELECTORAL_VOTES% is associated with a .38% decrease in the final vote spread. Advertising has no significant effect. We control for the September 27 probability of Obama winning the election, which is strongly associated with the final outcome. Heteroskedasticity-adjusted standard errors are displayed in brackets.

Variable	Coefficient
Obama Event Share Per Electoral Vote	1.26**
	[0.49]
McCain Event Share Per Electoral Vote	-0.38**
	[0.14]
Obama Advertising Share Per Electoral Vote	-0.70
	[0.98]
McCain Advertising Share Per Electoral Vote	0.28
	[1.33]
Initial Probability of Victory for Obama	0.23***
	[0.05]
_const	-6.72**
	[2.56]
Observations	15
R-squared	0.90

*significant at 10% **significant at 5% ***significant at 1%

Testing for robustness, we confirm the reasonableness of making a linear approximation for this function, and we find no interaction effect between Obama's events and McCain's events (i.e. the impact of a given event was independent of the number of events hosted by the opponent in that state). Additionally, we find no significant effect of allocating additional advertising spending to a particular state.

Unsurprisingly, we find that the September 27, 2008 Intrade price in a state is a very strong predictor of the eventual margin of victory. Figure 4.3 shows the relationship between the prediction market price and the final margin of victory.

To better comprehend the relative magnitudes of these metrics, consider a state with 10 electoral votes and an initial 40% probability of Obama winning. If no resources were allocated to this state, the final vote spread would be predicted as 40*0.23-6.72, or 2.48% in favor of Obama (with national-level fixed effects causing the shift in Obama's favor). To boost this spread by 1.25%, Obama could place 10% of his events in the state. Alternatively, a 5.4% increase in the initial probability would have the same effect. Since initial probabilities vary much more (by as much as 50%), the starting probability is the single largest factor in determining the final result. That is because the starting

Figure 4.3: Predictive Power of Electronic Prediction Market Price

This figure compares the final vote spread in favor of Barack Obama with the September 27 WTA probability for Obama in each state. There is a strong positive association.

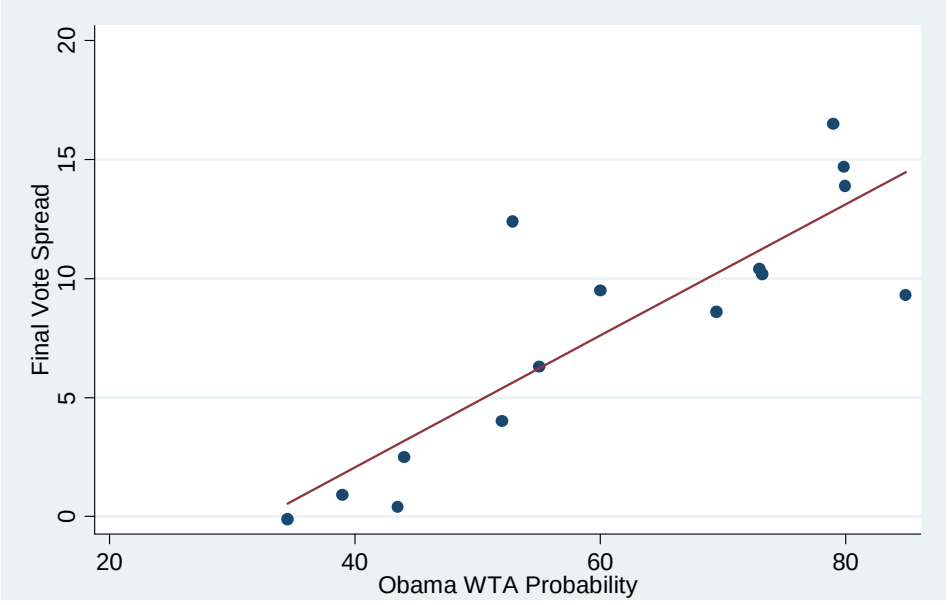

probability represents past effort and resources allocated to a state, as well as the bias of the state's median voter. Moreover, national-level contextual shifts likely perform a transformation on the initial probability of victory.

One intriguing result is the difference in impact between events hosted by McCain and those hosted by Obama; an increase in share of Obama's events caused nearly four times as much impact as the same increase in share of McCain's events. This result has several possible explanations. One hypothesis would be that Obama, who was able to generate large crowds and garner significant media attention throughout the campaign, held more effective events and the impact of resources on informational spread was greater for Obama's campaign than McCain's, or $\partial d_{ts,obama}/\partial o_{obama,s} > \partial d_{ts,mccain}/\partial o_{mccain,s}$ on average. Alternatively, viewing the election more holistically, average voters generally leaned towards Obama in this election. Thus, if ρ_{ts} is greater for Obama than McCain on average, then the same amount of information distributed by Obama could have greater impact; voters hearing news clips and watching speeches about positions with which they agree with more on average could reinforce their views and sway their vote.

The results of this chapter raise a number of interesting questions which beg further exploration. Despite the insignificant effect of advertising spending, candidates in practice exert a substantial amount of effort raising campaign funds to spend on advertising in key states. In the one-week period of our sample alone, McCain spent $9.5 million and Obama spent $16.3 million. If there is actually no effect of advertising spending on states, we would expect candidates to instead allocate their time towards having more campaign events, rather than spending time both fundraising and hosting

events. One hypothesis is that while local campaign events have a state-specific effect, advertisements have a national effect. If advertisements are rebroadcast across the country as part of a news cycle, their effect on individual voters may be national, as r_{ij} increases in all states rather than just the targeted state – and thus the candidate's decision for where to allocate advertising spending becomes inconsequential. Our model only includes relative impact on a specific state, and it does not observe national changes. A future direction of research may be to understand the national effect of advertising efforts and attempt to confirm this hypothesis.

As we noted in Equation 4.1, in addition to the two informational methods of appealing to voters we outlined, there is another way to increase expected agreement with the median voter: shifting views directly to correspond more with the electorate. Another future direction of research would be to understand the decision of whether to change political views and the impact of that choice.

Concluding Remarks

In this chapter, we utilize electronic prediction markets to measure the impact of candidate efforts to affect the election through strategic allocation of campaign resources. We find that there is a significant local impact of campaign events on the final margin of victory. Surprisingly, we find no local impact of allocation of advertising budget.

While the use of electronic prediction markets in this chapter provides us with meaningful results, there were several limitations. Rather than challenging our core premise that prediction markets are useful tools to understand elections, these are in fact exceptions that lend support to our argument. The limitations within this chapter are due

to incomplete markets, and they could be solved in the future by additional prediction markets. For example, one strong assumption (which we partially relaxed) was that prediction market prices did not incorporate expectations of future campaign resource allocation. To solve this limitation, the best methodology would be to create a new prediction market with a payoff tied to the amount of future campaign spending or number of campaign events in a given state. This would allow the impact of changes in expectations to be measured directly. We could similarly resolve several of the other assumptions and limitations with even more robust markets.

While the results of this chapter are narrower and based on stronger assumptions than the results in other chapters, the methodology provides an example of the potency of electronic prediction markets in understanding political decision-making. The major result of this chapter is that – while we have shown that political and economic context plays a large role in elections – candidates can indeed impact the election themselves by controlling informational distribution. While this effect is much smaller than conventional wisdom would hold, large-scale presidential campaigns do indeed play a significant role in the political system.

More broadly, this chapter represents an additional application of our electronic prediction market methodology, demonstrating its ability to dissect candidate actions and compare events in multiple states. The analysis we perform is only possible because of the core premise of our paper – that electronic prediction markets are efficient and unbiased.

Chapter V

Victory by Association:
Measuring Coattail Effects

In this chapter, we study the magnitude of "coattail effects" in the 2008 election, or the impact of the presidential election on congressional elections. We utilize data from electronic prediction markets to measure these effects. In order to eliminate endogeneity biases in our analysis, we measure coattails by using the occurrence of a major event as an instrument. We select days in which major events occurred (such as presidential debates) that affected the presidential election without any direct effect on local elections. Since the shifts that occur on these days in congressional markets are due exclusively to coattails, we then measure the impact the exogenous events had on the congressional elections. We find strong coattail effects in House elections and insignificant effects in the overall Senate race. We then apply our methodology to the Minnesota Senate election, where we find strikingly strong coattail effects. We discuss these results in the context of our simple model of voting preferences.

Introduction

In the chapter, we turn our methodology towards understanding local elections. While Chapter IV demonstrates that presidential candidates have some control over their electoral fate, Chapter III shows the decisive impact of economic and financial factors on the election – determinants that are outside a candidate's control. In addition to economic context, local elections are often shaped by another form of uncontrollable context: the national election. While resting partially on local issues and candidate personalities, local elections are often at least partially a referendum on national political figures. Local elections are heavily correlated with national results, and many elections – such as the 2008 election – have the same party sweep to victory in the House, the Senate, and the presidency.

Some of this correlation is due to contextual shifts, such as macroeconomic changes or wars abroad. Another source of correlation may be due demographic transition, leading to partisan realignment. A third part of the correlation comes from the presence of coattail effects, or the effect that a top-level candidate has on a local candidate. A popular presidential candidate can register new voters, drive partisan turnout and party-line voting, and change voter preferences on issues, leading to coattail effects that change the outcome of congressional elections.

Despite the intuitive nature of coattail effects, attempts at quantifying these effects have generally floundered. Since elections are correlated for a wide variety of reasons, measuring basic correlation between elections is insufficient to establish coattail effects. The methodological obstacles to measuring coattail effects were pointed out by Miller (1955), who noted the simultaneity of decisions taking place in elections. Mondak (1990) asserted that "a growing consensus holds that the presidential vote does exert

significant influence on congressional elections," though he added that these analyses were set back by the "long-recognized difficulties associated with measuring political coattails."

Despite the difficulty of measuring coattails, several attempts have been made. Kaplowitz (1970) proposed a metric called C-Correlation, which measures the correlation of local candidate results and presidential results after accounting for the number of swing voters in an election. While this demonstrates the degree of correlation between presidential and local politics, it does not disentangle which effects are due to presidential coattails and which, for instance, may be due to the economy's impact on Democratic candidates at both levels.

As an alternative methodology, in order to observe the change in coattail effects over time, Ferejohn and Calvert (1984) measured coattails as the impact of presidential preference on congressional elections after controlling for a wide range of economic variables. Similarly, Mondak (1990) measured the strength of coattail effects by observing congressional preferences after controlling for the level of economic concern and party affiliation of individuals within a single district. Since several factors play decisive roles in elections, accounting for these effects can substantially reduce false indications of coattails. Yet, unless these models control for all contextual factors within an election (including economic, social, international, and any other contextual changes that could shape an election), there will still be an endogeneity bias due to voters basing their decisions in both presidential and congressional elections on contextual changes.

Two primary obstacles exist in measuring coattail effects. First, there is a highly limited sample size: there is only one presidential election every four years from which to

draw data, and there are substantial contextual changes that occur over that time frame. To overcome the limited datasets, we use electronic prediction markets to assess the status of elections on a continual basis. A second obstacle is one of endogeneity: much of the correlation between candidates is due to agreement on issues and shifts in economic context, rather than coattail effects. In this chapter, we measure coattail effects within the 2008 election by using exogenous changes in the presidential election (that do not directly affect congressional elections) as an instrument. We find substantial coattail effects in the House and insignificant effects in the Senate.

The chapter is organized as follows: first, we describe a theoretical model for the existence of coattail effects; then, we propose an empirical methodology based on the electronic prediction market dataset; next, we analyze the results of this methodology; and finally, we provide concluding remarks and suggest areas for future research.

Modeling Coattail Effects

Despite the difficulties in measurement, there are intuitive reasons why coattail effects are likely to exist. Several institutional phenomena suggest a theoretical justification for coattail effects to play a large role in elections. First, voter registration and turnout are often driven by grassroots campaigns and excitement for a top-level candidate. If a presidential candidate is able to register millions of first-time voters, those voters are likely to vote in lower-level elections for the same party. Similarly, excitement about a top-level candidate can help drive party-line voting in some states, as voters that are ambivalent about lower-level elections choose out of simplicity and expediency to vote for a party generally rather than evaluating individual candidates. While these effects can help shift votes towards a popular presidential candidate's party, a more subtle

source of coattail effects shapes the median voter's preferences directly in lower-level elections. The median voter, typically a swing voter who was planning to vote and does not make a straight party-line vote, may change her lower-level preferences based on a change in her perception of a top-level candidate; her preference for a particular presidential candidate informs her decision of which lower-level candidate to support.

To gain a better understanding of the impact of coattail effects on the election, we will consider a congressional election that is simultaneous with a presidential election. Define candidate C as a congressional candidate in the same party as presidential candidate J. While the presidential candidate was modeled in Equation 2.3, we know that the candidate C can be simultaneously modeled as follows:

$X_{ic} = 1$ if $U_{dif\_c} \geq \overline{U}$ **(From Equation 2.1)**

$X_{ic} = 0$ if $-\overline{U} < U_{dif\_c} < \overline{U}$

$X_{ic} = -1$ if $U_{dif\_c} \leq -\overline{U}$

Where $U_{dif\_c} = E[U_{ic}]-E[U_{ic2}]$

X_{ic}: vote cast by voter i for candidate c. Not voting counts as 0 votes, and voting against counts as a negative vote for the candidate.
U_{ic}: utility voter i receives in the state of the world where candidate c wins
U_{ic2}: utility voter i receives in the state of the world where candidate c's opponent wins
$\overline{U}$: absolute utility difference threshold that causes voter i to cast a vote.

Additionally, we know:

$E(U_{ic}) = E(F(i_i, v_{ic}))$ **(From Equation 2.3)**
$i_i = i(c, t_i)$
$E(v_{ic})=v(c, r_{ic}, p_c, s_{ic})$

U_{ic}: utility voter i receives in the state of the world where candidate c wins
v_{ic}: vector of degree of agreement between voter i and candidate c's decisions
i_i: vector of weights of importance that a voter gives to particular issues
c: context of election

t_i: voter i's tastes
p_c: personal characteristics of candidate c
r_{ic}: historical correlation between voter i and candidate c's known beliefs
s_{ic}: signals of future agreement between voter i and candidate c

Under this model, a voter's tastes (t_i) and the election's context (c) are the same in both the presidential and congressional elections, and are not affected by either candidate. Similarly, the personal characteristics of a candidate (p_c) and the agreement between a voter and a candidate's revealed beliefs (r_{ic}) are specific to a particular candidate, and we assume that they are not affected across elections. Instead, we model coattails as an informational phenomenon, with signaled beliefs (s_{ic}) as the source of coattail effects between the elections.

Intuitively, the model suggests that coattail effects occur as voters use presidential candidates as partial proxies for local, less-known candidates. Beyond any information known about the candidate personally, the candidate's party affiliation is also visible to a voter – a strong signal of the candidate's future decisions if elected. Mondak (1990) finds that voters without much knowledge of a political race may use their views of the presidential candidate in the same party as a factor in their votes. Given the time cost of information collection and the voter's scarce budget of time, a typical voter spends only part of her time considering the election, and uses the presidential candidate as a proxy for the views of the local candidate. We formalize this view of election signaling as a simplified model:

$$s_{ic} = s(v_{ij}, w_{ic}) \qquad \textbf{(Equation 5.1)}$$

s_{ic}: signals of future agreement between voter i and candidate c
v_{ij}: vector of expected agreement between voter i and presidential candidate j's decisions

w_{ic}: relative strength of presidential views as a proxy for congressional candidate's views for voter i

The extent that the voter expects to agree with the presidential candidate's political decisions (v_{ij}) affects the voter's expectation of how much she will agree with the lower-level candidate (v_{ic}). The importance of the presidential candidate as a proxy is given a weight (w_{ic}) that shows the extent to which the voter relies on the presidential candidate as a proxy (which in turn is affected by the perceived correlation between the presidential candidate and the local candidate as well as the relative amount known about the views of each). Assuming a candidate's views are considered to be positively correlated with the presidential candidate of the same party, then $\partial s_{ic} / \partial v_{ij} \geq 0$, implying that $\partial U_{ic} / \partial v_{ij} \geq 0$. Thus, the probability that the voter votes for candidate c increases with the voter's expected degree of agreement with the presidential candidate (v_{ij}). Furthermore, the magnitude of impact of coattails on voter choice increases with weight, so $\partial^2 U_{ic} / \partial v_{ij} \partial w_{ic} \geq 0$.

This model confirms our intuition by outlining a mechanism for why, *ceteris paribus*, a candidate's popularity might be affected by a change in political popularity of the presidential candidate from the same party. Since a voter often knows more about a presidential candidate than about a local candidate, she uses the presidential candidate as a proxy for the local one – resulting in coattail effects.

The Problem with Traditional Measurement Approaches

For an observer trying to estimate the impact of coattail effects, many of the variables from Equations 2.3 and 5.1 are not directly observable; instead, what is observable is the median voter's decision on the presidential election, X_{ij} and their

decision in the lower-level election X_{ic}. Previous attempts at measurement, such as the C-Correlation approach, have effectively regressed lower-level election results directly on presidential election results, or X_{ic} on X_{ij}. Equation 2.3, however, demonstrates why this is problematic. The unobserved effect of context is a factor in both X_{ic} and X_{ij}, causing correlation without coattails.

To give a concrete example of this problem, in the 2008 election, negative events in the economy tended to shift voters' views in ways that were favorable to Democrats in general. There was substantial correlation between local elections and the presidential election, but much of this relationship was due to the favorable context for Democratic policies that were shared across candidates. The correlation in such a case should not be attributed to coattails, since it was economic changes – rather than changes in the presidential election – that shaped the association between the two elections.

A New Empirical Methodology

Despite the intuitive theoretical justifications for coattail effects, attempts to measure the effects empirically have been hindered by both causality questions and a limited data set. To overcome these obstacles, we propose an instrumental variable approach that uses data from electronic prediction markets. There are three steps to our methodology: first, we will broaden the dataset by turning to electronic prediction markets; second, we will select specific events where the presidential elections were exogenously affected without any direct effect on congressional elections; and finally, we will analyze the impact of these events on the congressional elections.

Step 1: Broadening the Data Set

In order to broaden the dataset, we use electronic prediction market data for the 2008 election from the Iowa Electronic Markets. We track daily closing prices for Democrats in the winner-take-all (WTA) market for the presidential election and the seat gain WTA market for the House and Senate elections from August 26, 2008 (the date the congressional markets opened) until November 4, 2008 (Election Day). In these markets, House and Senate prices are aggregated total probabilities of the Democrats winning seats in each respective chamber rather than looking specifically at particular districts. Figures 5.1 and 5.2 show the association between weekly changes in presidential markets and changes in the House and Senate races, respectively. As expected, the figures demonstrate a clear correlation, from which we intend to isolate the coattail effects.

Figure 5.1. House and Presidential Election Correlation

Figure 5.1 shows weekly changes in prediction market prices for Democrats winning seats in the House (on the y-axis) against weekly changes in price for Obama winning the election (on the x-axis). There is a general positive correlation between these two probabilities.

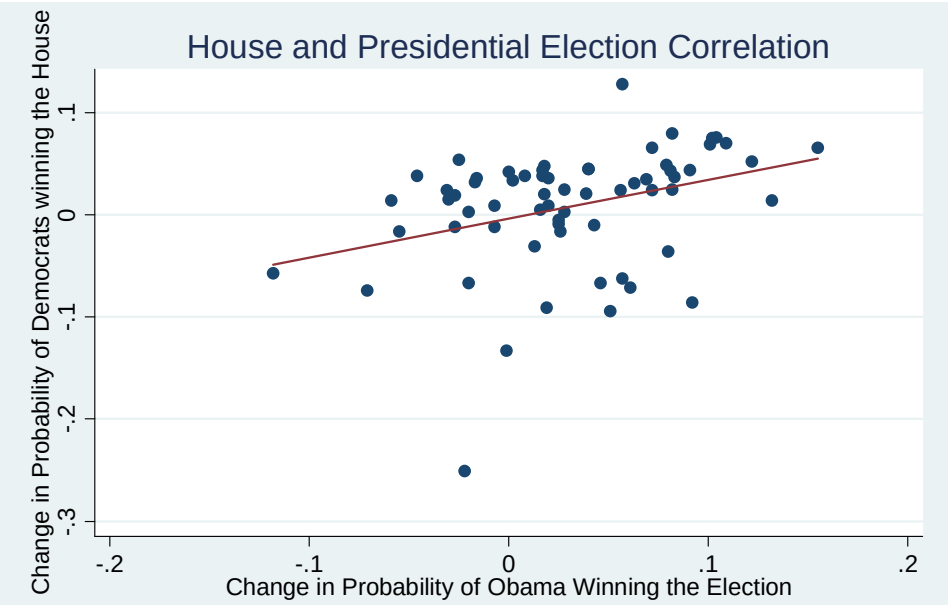

Figure 5.2. Senate and Presidential Election Correlation

Figure 5.2 plots weekly changes in prediction market prices for Democrats winning seats in the Senate (on the y-axis) against weekly changes in price for Obama winning the election. There is a general positive correlation.

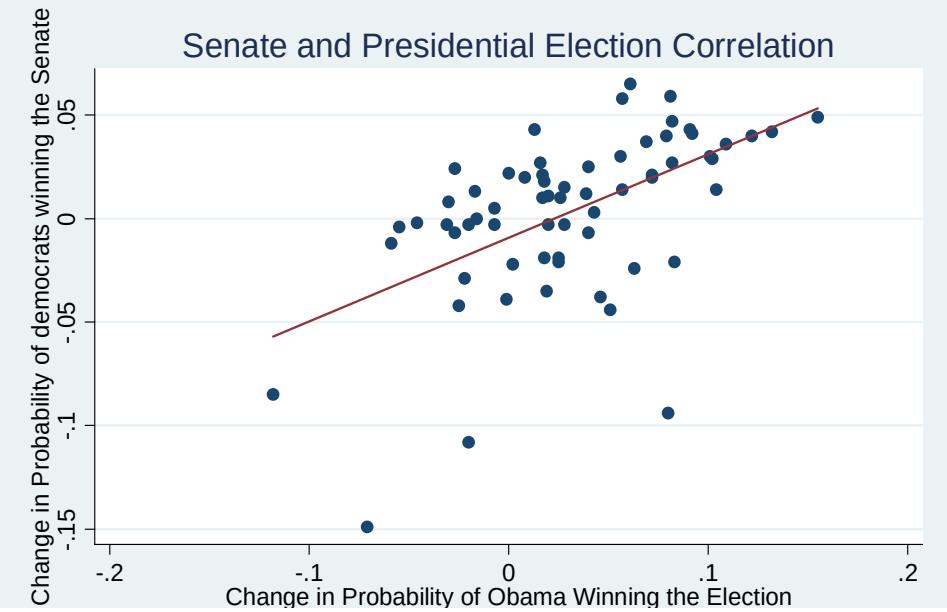

Table 5.1. Summary Statistics

Table 5.1 provides summary statistics for our dataset. The congressional Winner-Take-All (WTA) Price is the market probability that the Democrats win seats in a particular election.

(Variable)	(Obs)	(Mean)	(Std. Dev.)	(Min.)	(Max.)
Presidential WTA Price	46	0.693	0.106	.553	.903
Daily Change Presidential WTA Price	45	0.003	0.023	-.09	.054
Senate WTA Price	30	0.923	0.045	.806	.993
Daily Change Senate WTA Price	29	0.603	0.044	-.149	.059
House WTA Price	30	0.881	0.065	.764	.974
Daily Change House WTA Price	29	0.002	0.054	-.133	.069

Due to light trading and large bid-ask spreads, we use weekly price changes for the congressional prices instead of daily changes. However, if we were to compare weekly prices for each, we would unintentionally include correlation that does not result from coattails by expanding our time horizon too greatly (for example, economic changes over that time period may affect both if the time horizon is too long). Thus, throughout this chapter, we are generally comparing weekly congressional shifts with daily presidential shifts, assuming that any congressional price change that occurred in the preceding six days is in a random direction. We also remove from the dataset all periods in which there was no trading volume. Finally, we restrict the dataset to periods in which Obama gained vote share as a proxy for an election event being positive. Summary statistics are shown in Table 5.1.

Step 2: Choosing Events

In order to construct an instrument, we establish a binary variable dependent on whether or not there is a major shift that affects the presidential election without directly impacting lower-level elections. While some factors in the presidential election, such as economic context, are also major direct factors in lower-level elections, certain events only affect the presidential election directly. For example, a presidential debate – designed to influence public opinion about only the presidential candidates – would impact congressional outcomes exclusively through coattail effects. After analyzing major events in the election,[1] we select seven such events to set our binary variable equal to one:

[1] While a number of resources were used to select the largest events of the election, the Pew Research Center's "Top 25 Events of 2008 Election" was particularly useful. In choosing events, we selected events that were clearly exogenous to the election and were confined to a discrete time period. While there

- 8/23/08 - Obama selects Biden for VP
- 8/28/08 - Obama Acceptance Speech
- 8/29/08 - McCain selects Palin for VP
- 9/27/08 - Day after 1st debate
- 10/3/08 - Day after VP debate
- 10/8/08 - Day after 2nd debate
- 10/16/08 - Day after 3rd debate

To confirm that these events had a substantial effect on the presidential election, we compare prediction market price changes on days in which Obama gained to positive changes on days when there was a major event. As Figure 5.3 shows, the average gain in the presidential markets nearly doubles on the dates with major events.

Figure 5.3. Average Gains in the Presidential Election

Figure 5.3 displays the average daily gains for the Obama WTA contract in the electronic prediction markets on days where the closing price increased. The overall average daily gain is smaller than gains that take place on days with major events.

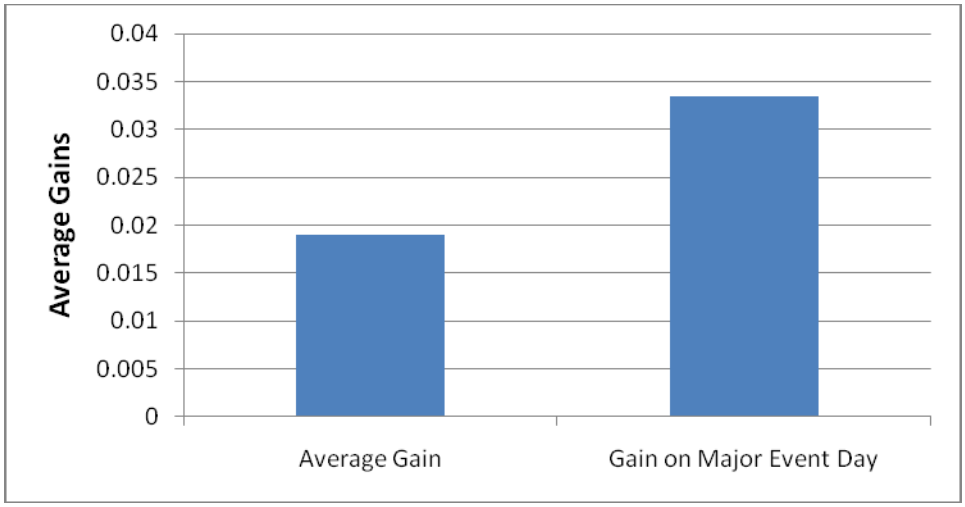

were several other significant events in the Presidential election (such as Obama travelling to Europe), we specifically selected for events with measurable short-term effects.

Step 3: Measuring Impact

Now that we have measured the impact of these events on the presidential elections, we turn to the congressional elections to gauge the impact that these major, exogenous events from the presidential election have on the congressional markets. Figure 5.4 shows a substantial effect of major presidential events on the congressional markets that increases average gains by a factor of eight in the Senate and a factor of four in the House.

To quantify this effect, we run a two-stage least squares regression using the presence of a major event as an instrument. As our first step, we regress changes in the probability of Obama winning on the presence of a major event on days where Obama

Figure 5.4. Average Gains in Congressional Markets

Figure 5.4 displays the average weekly gain in both the House and Senate markets on days when there was a gain in the presidential markets. Days with major events showed substantially larger average gains.

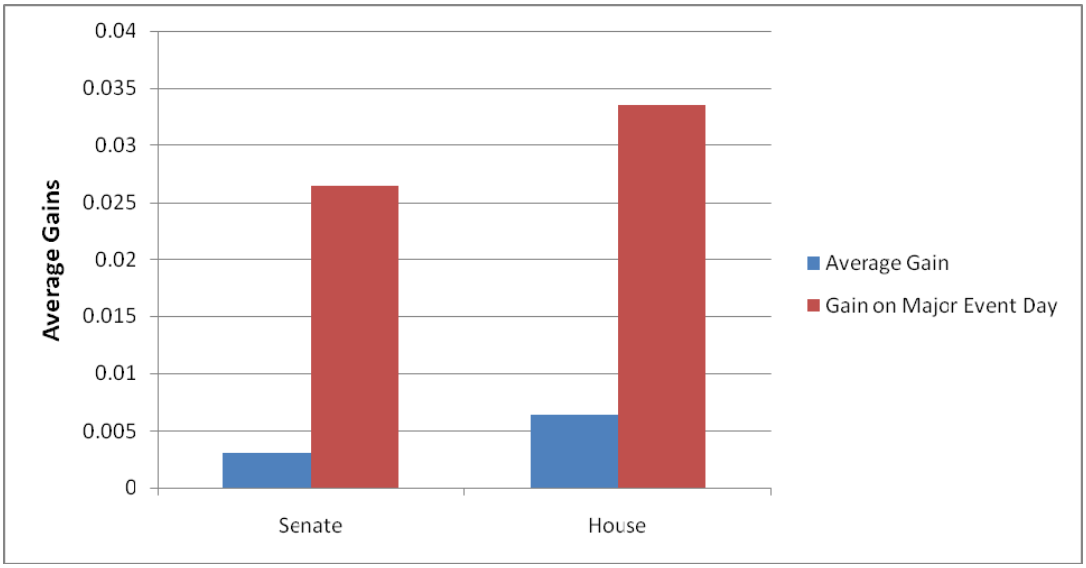

gained popularity in the election:

$$Change\_DemPres\_WTA = \alpha_0 + \alpha_1 * MajorEvent + u \quad \textbf{(Equation 5.2)}$$

Change_DemPres_WTA: daily price change for prediction market outcome of Obama winning presidency
Major_Event: binary variable equal to one on days with major presidential events
u: error term

We can now plug this result in to measure coattail effects. Equations 2.3 and 5.1 show that the probability that a particular local candidate wins the median voter's support is a function of candidate-specific effects (including the candidate likability and the agreement between the voter and the candidate's revealed beliefs), time-specific contextual effects, and changes in the popularity of the party's presidential candidate. In a naïve model regressing change in congressional probability on change in presidential probability directly, the error term would include both the candidate-specific effects and the time-specific contextual effects and would thus be correlated with the regressor. To overcome this problem with the naïve model, we instead use the predicted effect from Equation 5.2 as a regressor. Since we constructed *MajorEvents* to be uncorrelated with macroeconomic contextual changes and based exclusively on exogenous events in the campaign, its expected correlation with the error term is 0. Aggregating local elections into a metric of all ongoing congressional elections, this instrument enhances our model for a change in probability of the Democrats winning the election:

(Equation 5.3)
$$Change\_DemCongress\_WTA = \beta_0 + \beta_1 * Predicted\_Change\_DemPres\_WTA + \varepsilon$$

Where $Predicted\_Change\_DemPres\_WTA = \alpha_0 + \alpha_1 * MajorEvent$ (with α_0 and α_1 from Equation 5.2)

Change_DemCongress_WTA: weekly price change for prediction market outcome of Democrats winning seats Congress
Change_DemPres_WTA: daily price change for prediction market outcome of Obama winning presidency
ε: error term

Through the two-stage least squares outlined in Equation 5.2 and Equation 5.3, we observe the effect of an exogenous shift in presidential preferences on congressional prediction markets – yielding an unbiased estimate of coattail effects as β_1.

Results and Discussion

General Results and Discussion

Our results from the two-stage regression in Equations 5.2 and 5.3 are reflected in Table 5.2. The results display strong, significant coattail effects in the House election but weaker, insignificant results in the Senate. A 1% increase in Obama's likelihood of

Table 5.2: 2-Stage Least Squares Results

This table shows the results of Equation 5.3, demonstrating the predicted effect of exogenous changes in presidential market probabilities on House and Senate probabilities. The House demonstrates strong, significant coattail effects, while the effects in the Senate are minimal. Heteroskedasticity-adjusted standard errors are displayed in brackets.

Variable	(House)	(Senate)
Predicted_Change_Pres_WTA_Price	2.102**	0.764
	[0.901]	[0.587]
_const	-0.020	-0.005
	[0.021]	[0.015]
Observations	29	29

*significant at 10% **significant at 5% ***significant at 1%

winning is associated with a 2.1% increase in the probability of the Democrats winning the House, but it has an insignificant effect on the Senate. The coattail effect in the House election is quite strong, as an event occurring in the presidential election has a *larger* effect on candidates in the House than on the presidential candidates themselves.

The general model of coattail effects in Equation 5.1 helps explain the difference in magnitude of the effects in the House and Senate. One of the model's core results is that a voter's perception of the presidential candidate matters in congressional elections only to the extent that the presidential candidate's views are a useful proxy for the congressional candidate's views, or the magnitude of the variable w_{ic}. In an election where the median voter knows less about the candidate, coattail effects are expected to be stronger, since the party affiliation of the candidate carries more meaning. Typically, Senate candidates are better known than their House counterparts; since senators serve longer incumbencies, wield more personal power, spend more on their campaigns, and have broader constituencies, they generally capture a larger share of news coverage and political buzz. Thus, the median voter often knows more about the candidates in the Senate election than the candidates in the House election, and presidential views are more necessary as a proxy in the House election.

Amplification

One of the more surprising results of this analysis is the amplification of effects across elections. A one percent change in the probability of Obama winning is transformed into a greater than one percent change in the probability of Democrats winning in the House election. We propose two hypotheses for why amplification exists: geographic influences and the self-fulfillment of predicted success. In these markets,

changes in voting trends in different states are not equally important, as a vote shift in a non-swing state is not as meaningful as a vote in a swing state. The presidential swing states may vary from congressional swing states. Thus, a presidential announcement that "rallies the base" may have an amplified effect in states with close congressional elections but that are heavily leaning towards a particular party already in the presidential election.

An alternative hypothesis for amplification is that success is self-fulfilling, as the perception of being likely to win an election can increase the actual likelihood of winning. Campaign contributors, key politician endorsers, and even news editorial boards are often eager to support the winning side of an election and reap potential benefits of early support. A company's executives may choose to donate to a campaign they expect will win in order to gain potential political goodwill, and they must base their decisions on probability estimates made weeks or months before an election. Since these acts of support may directly shift the outcome of an election (and likely have larger effects in elections where candidates are lesser known), success may have inertia that amplifies a small probability change into strong gains.

Case Study: Minnesota

Turning from the general application of our model across all congressional and Senate elections in 2008, we now apply our model to one specific election: the 2008 Minnesota Senate election. Because of the perceived closeness of Minnesota's election, the Iowa Electronic Markets introduced a market with reasonable trading volume that measured the vote share in state's senatorial election. Table 5.3 shows the results of applying the instrumental variable regression from Equation 5.3 on Minnesota's results.

A 1% change in likelihood of Obama's victory is associated with a 5.865% increase in the expected vote share for Franken (the Democratic candidate in this election).

The magnitude of the coattail effect is strikingly large, especially for a Senate race. The difference in effect between this election and other Senate elections is likely due to the extreme closeness of this race – a race that was ultimately decided by just a few hundred heavily-contested ballots. Thus the impact of any perceived political shift would be heavily amplified in this ultra-close election. Furthermore, Minnesota was not a heavy swing state in the presidential election, and ultimately voted by a substantial margin for Obama. Thus, events in the presidential election that resounded with liberal voters may have had a stronger effect on the Senate election outcome than the general election probabilities.

Table 5.3: 2-Stage Least Squares Results in Minnesota

This table shows the results of Equation 5.3, demonstrating the predicted effect of exogenous changes in presidential market probabilities on the Minnesota Senate election. The election demonstrates strong, significant coattail effects. Heteroskedasticity-adjusted standard errors are displayed in brackets.

Variable	Coefficient
Predicted_Change_Pres_WTA_Price	5.865 *** [1.349]
_const	-0.057 [0.035]
Observations	29

*significant at 10% **significant at 5% ***significant at 1%

Limitations and Further Research

There are several limitations to these conclusions, and several questions are raised for further research. First, while this methodology could be generalized, we only used data from the 2008 elections in this analysis due to the limited history of electronic prediction markets for congressional elections. As more elections occur and the use of electronic prediction markets expands, a broader analysis could be performed on the impact of coattail effects on elections and the determinants of the magnitude of the effects. In particular, the 2008 election was unique in recent elections because of the relative certainty of the eventual outcome for Senate and House control, as the probabilities that Democrats would gain seats in each chamber were over 90% for much of the period evaluated. Thus, an analysis performed in a year when both the presidential and congressional elections are expected to be close may yield different results.

As prediction markets continue to mature, another potential expansion would be to observe coattail effects on a state-by-state level. While state-by-state election details from electronic prediction markets were limited in the 2008 election, further expansion in market availability and trading volume would enable a study on state-specific coattail effects and determinants of the magnitude of coattails.

Finally, several assumptions were made in choosing to use electronic prediction market prices as proxies for probabilities of actual outcomes. First, we followed our core assumption that prediction market prices are unbiased, reasonably accurate estimators. A more subtle assumption is that changes in prediction market prices are due to events impacting the election directly, rather than merely changing perceptions of the candidates themselves by market participants. For example, an eloquent speech by a candidate

would not only boost his chances in the election directly, it might also signal to prediction market participants that he will be similarly eloquent in the future – and traders will factor these future events into the market price. This effect would likely be particularly magnified early in an election, since market participants know less about the strategic and tactical competence of particular candidates. In order to mitigate this effect, we selected events for this study that occurred relatively late in the election and would likely have a primary effect of directly altering the election, such as the selection of a vice-presidential candidate. Furthermore, even if the market changes its perception of a candidate and incorporates information about future expected events because of one of the events we selected, as long as those future events also generate coattails they will also be incorporated into congressional prediction markets, leaving our estimate unbiased.

Concluding Remarks

Our results provide an intriguing insight into one of the major determinants of United States elections outcomes. Using electronic prediction market trading data, we find that coattail effects have a major impact on House elections, but a more limited impact on Senate elections. By using major exogenous shifts in the presidential election as an instrument, we isolate the coattail effects in the election and overcome causality concerns intrinsic to most past studies of empirical effects.

While this result generally confirms our intuitive expectations about coattail effects, the methodology complements the previous chapters in providing an intriguing new way to study major factors shaping elections. By assuming the market prices are robust, real-time indicators of the status of an election, our results show that – beyond

their predictive value – prediction markets can also be a useful tool in isolating empirical effects in elections and uncovering the determinants of political success.

Conclusion

In this paper, we have proposed a new methodology for understanding elections. The paper is based on a central assumption: electronic prediction markets are efficient and unbiased. By making this assumption – which is supported by both empirical and theoretical work – we are able to gain a deeper understanding of political events than would otherwise be possible.

In Chapter III, we found that the economic crisis played a major role in the 2008 election to Barack Obama's advantage, as we isolated a causal impact of equity market decreases, oil market increases, and 5-year T-Bill price increases in Obama's favor. Furthermore, we found that Obama's victory significantly boosted equity markets on Election Day.

In Chapter IV, we determined the impact of campaign spending and campaign events on the election. We used electronic prediction market prices to control for the starting point of an election, and we found that campaign events had a significant effect while advertising had no significant local effect.

In Chapter V, we measured the size of coattail effects in the 2008 election. By using electronic prediction market prices to eliminate an endogeneity bias, we found strong coattails in the House of Representatives and weak coattails in the Senate, confirming our predictions based on an intuitive model of coattails as an informational signal.

More important than the results we found, however, is the utility of the methodology. While many of our findings are intuitive, our methodology is a novel approach to solving a long-standing set of empirical challenges, and our dataset will only

grow more robust as electronic prediction markets continue to grow. Event studies, traditionally performed primarily for events in financial markets, are now possible on elections. Using our approach, several natural experiments take place within an election, and the event can be thoroughly understood. Further studies could use this methodology to understand major determinants of past elections (markets have existed since 1988), elections in other countries, and commonalities across countries. Furthermore, as more and more elections accumulate in which prediction markets are available (and markets become more accurate as more and more traders participate), cross-election models can be formulated, and Fair's model can be expanded to hundreds of continuous variables without concern of overparameterization. Understanding elections provides a glimpse at both the heart of democracy and the utility-maximizing political preferences of individuals; the improved ability to dissect elections and understand the factors contributing to their outcomes can yield deeply rewarding and useful information to both political economics and political science.

References

1. Berg, J. et al., "Results from a Dozen Years of Election Futures Markets Research." Working Paper (2001).

2. Berg, J. et al., "Accuracy and Forecast Standard Error of Prediction Markets." Working Paper (July 2003).

3. Erikson, Robert et al. "Was the 2000 Election Predictable?" *Political Science and Politics,* Vol. 34, No. 4 (Dec., 2001), pp. 815-819

4. Fair, Ray. The Effect of Economic Events on Votes for President. *The Review of Economics and Statistics,* Vol. 60 Number 2 (May 1978), Pages 159-173.

5. Fair, Ray. "Econometrics and Presidential Elections." *Journal of Economic Perspectives*. Volume 10, Pgs 89-102. Summer 1996.

6. Fair, Ray. *Predicting Presidential Elections and Other Things*. Stanford University Press (2002).

7. Ferejohn, John and Randall Calvert. "Presidential Coattails in Historical Perspective." *American Journal of Political Science*, Vol. 28, No. 1 (Feb., 1984), pp. 127-146.

8. Frey, B. and Schneider, F. "Economic and Personality Determinants of Presidential Popularity." *Empirical Economics,* Vol. 3, No. 2 (June, 1978), pp. 79-89

9. Kaplowitz, Stan. "Using Aggregate Data to Measure Presidential Coat-Tails Effects." Working Paper (1970).

10. Levitt, Steven. "Using Repeat Challengers to Estimate the Effect of Campaign Spending on Election Outcomes in the U.S. House." *The Journal of Political Economy.* Vol. 102, No. 4 (Aug., 1994), pp. 777-798.

11. Lewis-Beck, M. "Economic voting: an introduction." *Electoral Studies,* Volume 19, Issue 2 – 3 (2000), Pages 113-121

12. Miller, Warren. "Presidential Coattails: A Study in Political Myth and Methodology." *The Public Opinion Quarterly*, Vol. 19, No. 4 (Winter 1955), pp. 353-368

13. Mondak, Jeffery. "Determinants of Coattail Voting." *Political Behavior,* Vol. 12, No. 3 (Sep., 1990), pp. 265-288.

14. Mondak, Jeffery and Carl McCurley. "Cognitive Efficiency and the Congressional Vote: The Psychology of Coattail Voting." *Political Research Quarterly*, Vol. 47, No. 1 (1994), pp. 151-175.

15. *New York Times*. President Map. http://elections.nytimes.com/2008/results/president/map.html. Accessed February 2009.

16. Pew Research Center. "Top Events of Campaign 2008."November 6, 2008. Available at http://pewresearch.org/pubs/1025/election-news-interest

17. Shaw, Daron. "The Effect of TV Ads and Candidate Appearances on Statewide Presidential Votes, 1988-96." *The American Political Science Review*, Vol. 93, No. 2 (Jun., 1999), pp. 345-361

18. Snowberg, Erik et al. "Partisan Impacts on the Economy: Evidence from Prediction Markets and Close Elections." *The Quarterly Journal of Economics*, Vol. 122, No. 2 (2007), pp. 807-829.

19. *Washington Post.* Campaign Tracker. http://projects.washingtonpost.com/2008-presidential-candidates/tracker/candidates/barack-obama/states/oh/. Accessed in February 2009.

20. Wisconsin Ad Project. "Pres. TV advertising spending continues to grow; Over $28 million spent from September 28-October 4." University of Wisconsin. Available at http://wiscadproject.wisc.edu/wiscads_release_100808.pdf.

21. Wolfers, Justin and Eric Zitzewitz. "Interpreting Prediction Market Prices as Probabilities." *CEPR Discussion Paper No 5676* (May 2006).

22. Wolfers, Justin. "Best Bet for Next President: Prediction Markets." *Wall Street Journal.* December 31, 2007.