Volatility in Prediction Markets: A Measure of Information Flow in Political Campaigns

Carlos M. Carvalho[*]
&
Jill Rickershauser[†]

The University of Chicago and American University

May 23, 2008

Chapter X of

The Handbook of Applied Bayesian Analysis

Eds: Tony O'Hagan & Mike West
Oxford University Press

Abstract

In each U.S. presidential campaign, anecdotes that seem to summarize the election emerge, such as Bill Clinton's campaign slogan, "It's the economy, stupid," in 1992 and John Kennedy's performance in the televised debate against Richard Nixon in 1960. However, some stories may seem to epitomize the campaign when they actually did not alter the outcome of the election or even relate to the underlying factors that did influence the outcome. How can we adjudicate between the events that mattered and those that did not? Building on work in the market microstructure literature in economics, this paper builds a measure of information flow based on the returns and volume of the 'Bush wins the popular vote in 2004' futures contract on the Tradesports/Intrade prediction market. This measure allow us to associate a particular event to an information level, providing a direct way to evaluate its impact in the election. In the 2004 presidential race, our findings show that the televised debates, Kerry's acceptance speech at the Democratic convention, and several national security–related stories such as the report that explosives vanished in Iraq under the U.S.'s watch, the CBS story about Bush's National Guard service and the subsequent retraction, and the release of the bin Laden tape a few days before the election increased the information flow. Contrary to popular accounts of the election, we find that ads attacking Kerry's military service aired by the Swift Boat Veterans for Truth in August contributed only a limited amount of information to the campaign.

Keywords: Prediction Markets, Stochastic Volatility, Markov-switching, Political Campaigns, Campaign Effects, Information Flow.

[*]Graduate School of Business, The University of Chicago, Chicago, IL 60637
[†]Department of Government, American University, Washington, DC 20016

1 Introduction

When asked to describe the pivotal moments of the 2004 campaign, Mark Mellman, senior strategist for John Kerry's campaign, first noted that analysts should be careful to:

> *"Avoid what psychologists call fundamental attribution error. Fundamental attribution error consists, as some of you know, of overweighting the significance, the importance of individuals, of personalities, of events, and underweighting the significance, the salience of the structure of the situation of the underlying circumstances."*

He went on to outline the factors he believed to have been the most important in determining the vote in the 2004 election: the decline of the Democratic plurality among voters, culture as the primary division in American politics as opposed to class, and Bush's incumbency advantage. While he did list a number of events that the Kerry campaign viewed as important, his central message was that the campaign did the best it could and very well given the political situation in 2004. Matthew Dowd, a Bush strategist, also viewed the campaign as a series of relatively minor ups and downs with a largely expected outcome. Yet when most people think back to the 2004 campaign, some events seem pivotal, such as the Swift Boat Veterans for Truth (SBVT) ads in August and Kerry's comment that he voted for spending $87 billion before he voted against spending $87 billion for emergency funds for troops and reconstruction in Iraq. Did those two events matter much in the final vote tally, or did they not really change the campaign as much as we might think? Analysts look back at campaigns and identify defining moments but, as Popkin (1991) notes, those anecdotes can be simply the story that "fits" our overall impression of the campaign as opposed to the events that actually affected the election outcome. How can we tell the difference between anecdotal explanations and the events that were actually important? By identifying which campaign events did, in fact, provide new information in a given election year, we can better understand the "normal" dynamics of campaigns.

In order to investigate what kind of information and events have which kinds of effects on expectations' of a candidate's success, we need a measure that reacts quickly to new information. Opinion polls give us a somewhat nuanced view of public opinion during campaigns, but there is a practical problem with surveys: they are taken over several days. A poll taken the day before, the day of, and the day after a presidential debate does measure the effect of a candidate's performance in the debate on vote choice but it also registers the effects of events (e.g. candidate visits or speeches) that occur before and after the debate. If something shocking happened in a debate (would that debate–watchers could be so lucky!), the results from that standard three–day poll would partially reflect it, but we could only surmise what change in candidate preference occurred because of that gaffe and what change was caused by other events.

Prediction markets, like the Iowa Electronic Markets (IEM) and Intrade (previously called Tradesports), offer an attractive alternative to polls or academic panel studies to answer this kind of question because they allow us to clearly analyze people's expectations of political outcomes and the fluctuation in those expectations given new information. The pricing of futures contracts in these markets provides a continuous measure of the probability of a candidate winning the election. The immediacy of online futures markets lets us see exactly how information is incorporated, how quickly, and how expectations about political events stabilize.

This paper uses data from futures on 'Bush wins the popular vote in 2004', or the traded probability, of Bush winning the election, to build a measure of information flow. Our approach builds on theoretical developments in the economics literature, especially the market microstructure literature, that view the variation of asset prices as being driven by the arrival of new information and the process that incorporates that information into market prices. We extend a theoretically justified model that combines returns and volume in creating a measure of information flow to include two regimes of variation. This provides a novel approach to differentiate between the "normal" campaign occurrences and the noteworthy, important events of particular campaigns.

While the Swift Boat ads have become central to popular accounts of the 2004 presidential race, it appears that Kerry's nomination speech and the debates, particularly the third one, imparted much more information. Other events that were particularly important in the campaign were the CBS story about Bush's National Guard service and the subsequent retraction, the report that explosives went missing in Iraq in late October, and the release of the bin Laden tape

a few days before the election.

2 Political Prediction Markets

Prediction markets, which are also called "information markets," "idea futures," or "event futures," pay out returns based on whether or not a specified event occurs. Winner-take-all futures pay $1 if event X occurs and $0 if it does not. At one point during the 2004 campaign, a person could buy a 'Bush wins the popular vote' contract for $0.55. If that person held onto that contract until November 2, then she would receive $1, making a profit of $0.45 cents per contract. Investors can resell or buy contracts at any point. For example, she could buy a Bush contract at $0.55 and then sell it at $0.58, for a profit of $0.03 per contract. This paper examines winner-take-all futures markets, though there are several other types of contracts in prediction markets.

Before proceeding to a discussion of information flow in political prediction markets and the evidence about campaign effects in the 2004 election, a few theoretical points about these kinds of markets must be reviewed. In particular, we discuss why prediction markets are able to aggregate information such that the prices of future contracts are reflective of the event's actual probability of occurring.

How do markets aggregate information? When a group of people are asked how many marbles are in a given large jar, very few of them will choose the correct answer. However, the average of the guesses is always a close estimate of the actual number. Similarly, few people can estimate the number of yards between point A and point B, but together the group "finds" the answer. This occurs even without discussion; opinions are polled privately and then aggregated to determine the distance. This is an intuitive explanation of how a prediction market works. The theoretical explanation rests on the efficient market hypothesis and no arbitrage arguments (Wolfers and Zitzewitz, 2004).

Prediction markets are likely to produce quality information because they provide: (i) incentives to seek information; (ii) incentives for truthful information revelation and (iii) an algorithm for aggregating diverse opinions (Wolfers and Zitzewitz, 2004). These markets force people to *"put their money where their mouths are,"* so their expressed beliefs about a candidate's chances of success are not just cheap talk. It is important to note that the average belief of traders exhibited in the futures price is actually a weighted average, where people with more confidence in their assessment of a candidate's success trade more so that their opinions count more in the price equilibrium. This can be viewed as an advantage over traditional polls: equilibrium prices take into account the strength of people's beliefs about the candidate's chance of electoral success.

Without discussing a candidate's chances, the information held by people all over the country will be combined so that the market estimate of the candidate's probability of winning is accurate. That is, a person who lives in Ohio might trade based on his impression of voters in his town and a person in North Carolina might trade on her impression of the local newspaper coverage. When added together with traders from around the country (and around the world), all the information is combined to produce the accurate value of the contract.

2.1 Bounded Rationality

The economic theory behind prediction markets suggests that the prices of assets are accurate because people rationally and correctly understand and incorporate information into the market prices. However, behavioral research shows that people are limited information processors and are subject to biases due to their preferences. This is not a major problem because *"all agents need not be informed for a market to be efficient... survey evidence documenting biased expectations does not necessarily imply that market prices will be similar biased "* (Sauer, 1998). A market is efficient in the sense that it fully reflects all public information because of marginal traders or market makers. In a market of sufficient size and with large enough payoffs to counter transaction costs, arbitrage in prices will occur until all traders, specialists as well as informed and uninformed investors, agree on the appropriate price (Andersen, 1996). At this point in

time Tradesports represents the largest political prediction market and appear to have enough participation to alleviate the bias problem.

It is hard to deny that most traders are biased much of the time. Berg and Rietz (2006) notes that *"traders are biased and mistake prone"* and Forsythe et al. (1999) concur that *"traders frequently leave money on the table through violations of arbitrage restrictions."* These conclusions are documented by surveys administered by the Iowa Electronic Market to its traders. The surveys have consistently shown that traders are not in any way representative of the voting population. In particular, traders on the IEM are younger, more educated, more wealthy, and "more male" than the U.S. population and the voting population. Due to the incentives to seek information when trading, they are also probably more informed than the voting population.

Traders in the IEM tend to hold portfolios that are biased towards their preferred candidates (Forsythe et al., 1992). Forsythe et al. (1998) show that investors in the 1993 Canadian elections futures markets held unbalanced portfolios in ways consistent with their preferences. That is, supporters of a party were much more likely to hold more futures of that party than was the market as a whole. Forsythe et al. (1999) demonstrate that this bias also exists in the U.S. presidential prediction markets. The trading population of Tradesports is likely somewhat different than the trading population of the IEM due to its larger size and commercial aspect. However, there is still reason to suspect that individual bettors in Tradesports markets are also biased.

These types of biases have been described as "wish fulfillment" in both the political science literature and the psychological literature. Granberg and Brent (1983) and Uhlaner and Grofman (1986) document the persistence and ubiquity of this phenomenon. Supporters of Carter in 1980 were much more likely to predict that Carter would win than were Reagan supporters (though a few of each candidate's supporters were less biased in their predictions). Whether due to the "false consensus effect," where supporters of a particular candidate see themselves as more representative of other voters than they actually are, or the "assimilation-contrast effect," where supporters of a candidate interpret news about that candidate much more positively (Uhlaner and Grofman, 1986), most traders in the IEM exhibit this wishful thinking as evidenced in their portfolio holdings.

2.2 Prediction Market Efficiency

Demonstrations of these biases does not necessarily mean that the prices of candidate futures are incorrect. While there is some evidence that most traders in the IEM are biased in their trading decisions (or at least likely to make mistakes) and it is likely that many, if not most, of the traders in Tradesports markets are similarly behaved, markets are nonetheless accurate because of marginal traders. Though this paper specifically focuses on the effects of information on futures prices as opposed to the markets' predictive power, this section will briefly review the evidence of the accuracy of prices in prediction markets in order to show that analysis of prediction markets' prices is worthwhile.

Many studies of the IEM show that prediction markets are as accurate as or more accurate than opinion polls, both in the week or days before the election and at longer time horizons. Berg et al. (2003) note that since 1988, the IEM political prediction markets (conducted in U.S. elections as well as in other countries such as Canada, Austria, and Turkey) have consistently outperformed polls and *"in a few cases (the 1988 and 1992 U.S. Presidential Elections) the market dramatically outperformed polls."* Not surprisingly, markets with a higher volume of contracts traded and markets with fewer candidates or parties are more successful. U.S. presidential election markets are the most accurate, though the overall market accuracy is also impressive: the average market error in vote-share contracts[1] was 1.49% or 1.58% (depending on the exact measure used) whereas the average poll error was 1.91%. It is more difficult to judge the true probability of an event occurring, but in markets where a winner-take-all contract and a vote-share contract are both traded, evidence suggests that the prices of the two contracts move together such that the probability expressed in the winner-take-all contract is similarly accurate. The ability to predict the winner is intellectually interesting[2], though the accuracy of

[1] A vote-share contract pays out \$1* the percent of the vote won by the candidate or party.

[2] It is also potentially financially rewarding to predict the outcome of elections (see Hartog and Monroe, 2008).

the election markets has been studied so extensively in part because other applications of idea futures markets rely on the markets' ability to accurately forecast events.

If surveys of traders and findings from political science and psychology all show that people invested in the market are biased, how are market predictions themselves accurate? The apparent paradox is resolved by the "marginal trader" (e.g. Forsythe et al., 1992, 1999). As Forsythe et al. (1998) explain, *"as long as there are some traders relatively free of such wish fulfillment biases and with deep enough pockets, they will take advantage of the biases of other traders and in the process bring prices to levels consistent with unbiased expectations."* These less-biased investors serve as "market makers," which is to say that they will continue to trade until the price of each futures contract is correct.

Oliven and Rietz (2004) show that the average trader leaves money on the table, either due to bias or mistakes, but that the marginal trader takes advantage of those arbitrage opportunities and adjusts the prices. Other work has shown that market makers, the marginal traders, whose portfolios and trades exhibit no bias, invest much more money than the average trader and are more active in trading. As a result, marginal traders are able to drive prices to efficient levels while profiting from the mistakes of more error-prone traders (Forsythe et al., 1999).

Marginal traders' behavior explains the accuracy of prediction markets, so our starting premise is correct: these markets are worth studying especially in regards to campaign and information effects. While it is true that the equilibrium price in a market with positive information and transaction costs does not summarize some small amount of the information, that fraction is negligible (Grossman and Stiglitz, 1980). These findings suggest that prices of futures are, in fact, reasonable assessment of the "true" probabilities of a candidate's success given that all information is rationally incorporated into prices, even if individual traders are not always (or even usually) rational.

3 Volatility, Trading Volume and Information Flow

Studying the changes in the probability of a candidate winning an election is as important as studying the probability itself. In the context of prediction markets, this means that price movements and trading volume must also be considered in order to understand the dynamics of the campaign. In their study of presidential approval, Gronke and Brehm (2002) write:

> *"No one would argue that a statistical distribution can be described solely by its central tendency; nor should our exploration of the patterns and causes of presidential approval look only at the mean. As important in many circumstances is the frequency and sharpness of shifts in public sentiment-volatility."*

Similarly, if we are interested in the question of what information during campaigns is seen as important or consequential (or at least potentially consequential), the movement of futures prices matters. There may be a significant amount of action in a futures price when new information is acquired, though ultimately the price may settle back to the original equilibrium. An analysis of the mean price over one or several days would hide that effect of political information. Therefore, such an analysis would underestimate the perceived news content of specific campaign events.

Prices remain relatively stable when investors learn no new information about the underlying asset. However, when new information such as earnings reports does become public, stock and futures prices can move more dramatically as traders incorporate the arrival of information. This results in prices revealing the full information content through a sequence of trades. This phenomenon is well studied, with a vast literature in economists and finance focusing on the effects of earnings reports and news releases on volatility and volume (see Jennings and Starks, 1985; Brown and Hartzell, 2001; Patell and Wolfson, 1984; Chan, 2003; Dubinsky and Johannes, 2006) as well as the relationship between prices, volatility and volume (e.g. Gallant et al., 1992).

In analyzing the effects of politics on financial markets, Bernhard and Leblang (2006) pay careful attention to the volatility of asset prices as a measure of uncertainty about a political outcome. Their particular contribution is to highlight the fact that information affects markets conditional on prior expectations, so simple dummy variables in a time period in which some

event occurs can miss the actual impact of the event. When they control for expectations about the 2000 presidential election (as measured by poll numbers), they find that exit poll results from battleground states did affect the Standard and Poors and NASDAQ futures and the US dollar-Japanese yen exchange rate. Similarly, this paper tries to identify unexpected events that affected the electoral chances of the candidates in 2004, above and beyond the long-term, underlying electoral landscape.

Volatility can be conceptualized as a measure of the information flow into a market. In the market microstructure literature in economics, the mixture of distributions hypothesis *"posits a joint dependence of returns and volume on an underlying latent event or information flow variable"* (Andersen, 1996). While the theory describes the *"informativeness of market prices, the presence of liquidity traders in the market and the manner in which news is disseminated"* as factors influencing both volatility and volume, it *"points to the rate of information arrival to the market as the primary variable of interest"* (Fleming et al., 2006). The volatility, or the variance of the prices, is *"primarily [caused] by the arrival of new information and the process that incorporates this information into market prices"* (Andersen, 1996).

Information flows into the market and affects both the prices of the assets and the volume of the assets that are traded. When this trading occurs, the price of the asset reaches a new equilibrium state. Andersen (1996) explains:

> *"An important finding is that over the course of a (short) period, the sequence of trades reveals the pricing implications of the private signals and subsequently–until new private information arrives–all market participants agree on the value of the traded asset. Thus, private information arrivals induce a dynamic learning process that results in prices fully revealing the content of the private information through the sequence of trades and transaction prices."*

It is new information that starts this process. In the context of political prediction markets, when new information about the campaign–whether polling numbers, a larger ad buy in a battleground state, or an endorsement–becomes known, the price of the candidate's futures contract adjusts accordingly through the increased participation of market makers. Volatility is low despite some new information when rational anticipation of that event was already incorporated into the price. In contrast, high volatility occurs when unexpected information enters the market.

This price adjustment process, as evidenced in the estimated volatility, can be used to characterize information and events. Jennings and Starks (1985) study the effects of high-content earnings announcements as compared with low-content announcements and find that high-content ones induce more volatility for longer periods of time. Studies like that one take the information content as fixed and study the effects on the market, whereas this paper looks at the effects on the market and attempts to make statements about the information content contained in different campaign events. This is done by first estimating a measure of information flow (or volatility) using returns and volume. This measure is then used to evaluate the informativeness of a campaign event. While this approach presents some hurdles to making assertions about causality, it provides a novel and theoretically sound way to identify consequential events, therefore enhancing our understanding of the role of information, campaign events, and other factors on election outcomes.

In this paper we work with a extended version of Andersen (1996), where the underlying information flow process is determined by a Markov-switching stochastic volatility model (Carvalho and Lopes, 2007; So et al., 1998) that allows for occasional shifts in the parameter determining its level. The incorporation of structural changes in the model solves the problem of overestimation of the persistency of volatility in high-frequency data (Lamoureux and Lastrapes, 1990) while providing a straightforward way to identify clusters of high and low information flow in the market. In Andersen's model, returns and trading volume follow a mutually and serially independent bivariate distribution, conditional on the number of information arrivals, where information arrivals are modeled as a stochastic process. This is an empirical representation of a well defined market microstructure model in which both informed and liquidity traders (traders that do not react to news) participate in determining prices in the market and where the volatility, or information flow, process is inferred simultaneously from returns and trading volume.

In detail, let R_t represent the logarithm of returns, V_t a measure of trading volume and I_t the flow of information at time t. The model takes the following form:

$$\begin{aligned} R_t &\sim N(0, I_t) && (1)\\ V_t &\sim Po(m_0 + m_1 I_t) && (2)\\ log(I_t) &= \alpha_{S_t} + \beta log(I_{t-1}) + \epsilon_t && (3)\\ S_t &\sim \mathcal{MS}(P) && (4) \end{aligned}$$

where $\epsilon_t \sim N(0, \tau^2)$ and $\mathcal{MS}(P)$ denotes a two-stage first order Markov-switching process with transition probabilities defined by $Pr(S_t = i|S_{t-1} = j) = p_{ij}$, for $(i, j) = 1, 2$.

In the model, returns are conditionally normal with variances that reflect the intensity of information arrivals that also characterize how strongly volume fluctuates in response to news. In this context, m_0 represents the participation of liquidity traders in the market whereas the proportion of trades by "informed" traders (or market makers, see Section 2.2) is defined by m_1. Moreover, the information flow I_t (also refereed to as volatility) follows a conditionally stationary ($|\beta| < 1$) log AR(1) process where the level α_{S_t} is determined through a two-stage Markovian process. We will refer to S_t as the the level of information (or information state), meaning that when S_t equals 1 there is a reduced level of news arrival. This implies that $\alpha_1 < \alpha_2$ as these parameters define the unconditional mean of the information flow process. It is important to point out that the choice of two stages for S_t has, on the one hand, an applied motivation as we are looking for a way to segregate between periods of low and high information. This device is used by Carvalho and Lopes (2007) in flagging currency crises in emerging markets. On the other hand, evidence in the literature suggests that at least two stages are necessary to properly characterize the volatility of assets (So et al., 1998; Lopes and Carvalho, 2007; Eraker et al., 2003)[3].

Equations (1-4) represent a complex, bivariate dynamic non-linear system with two latent (non-observable) states, the information flow I_t and the information state S_t. Using standard priors, posterior estimation of this model is carried by a customized Markov chain Monte Carlo algorithm described in appendix A.

We implement the described model and construct a measure of information flow for the 2004 presidential race. The data were acquired from Tradesports. They include all transactions made on the Tradesports 'Bush wins the popular vote' contract from July 1 to Election Day. The transactions were grouped into 3 hour time periods. This choice reflects our hope to work with the finest level of resolution while guaranteing that enough transactions occurred in the period. Observations of the price of the contract at the close of the three hour period and the volume traded in that period were recorded at 3am, 6am, 9am, noon, 3pm, 6pm, 9pm, and 11:59pm. Following the standard in the literature (Andersen, 1996), volume was transformed to the log scale and rounded to the closest integer due to the Poisson model specification. Futures prices, returns and volume are presented in Figure 1 and 2. Finally, we catalog events from the Drudge Report (`www.drudgereport.com`) in the same 3 hour window. We use the Drudge Report for several reasons. First, to investigate whether the market reacts to campaign events, we need a "time-stamp" on a news story. Traditionally, media studies have "timed" events by using the day that the story was on the front page of a major newspaper such as the *New York Times*. Theoretically, we expect the market will react much more quickly, so we need a more fine-grained measure. The Drudge Report provides that; it is constantly updated and is archived every 15 minutes. Second, the Drudge Report is a widely-read political gossip source. In a recent book (Halperin and Harris, 2006), two influential political reporters write:

> *"It is a guarantee that most of the reporters, editors, producers, and talk show bookers who serve up the daily national buffet of news recently have checked out his eponymous website, and that www.drudgereport.com is bookmarked on their computers. That is one reason Drudge is the single most influential purveyor of information about American politics."*

[3]We did assess the performance of a model with 3 regimes and due to the similarity between the levels (α) of of two states we are confident that the model with 2 regimes is sufficient and appropriate for what we are trying to accomplish here.

People do use this as a source of information, and that information is on the "early" side–Drudge routinely reports stories that the major media news outlets have not yet reported or vetted. The most famous example is the Monica Lewinsky story; the Drudge Report was the first place to break most of the news, and covered the story for weeks before the major news outlets. Because it is an agenda-setter with early information and has precise time stamps, we feel that this is the best available source[4].

4 The 2004 Presidential Election

After fitting our model, we can initially conclude that the flow of information during the campaign was not static. Estimates of the states I_t and S_t are presented in Figure 3 and a posterior summary of all parameters appears in Table 1. Volatility rose significantly at many points during the campaign, though it remained low for much (but not all) of July, August, and September. The duration of increased volatility also varied. Spikes in volatility were often accompanied by longer stretches of a state of high volatility, indicating greater uncertainty about the effect of new information. However, on many occasions, the volatility increased to a high state and then quickly fell back to the normal level, indicating that the effects of some new information were easily understood and incorporated into the probability of Bush's electoral success. This is also reflected in the estimates of p_{11} and p_{22} given that $p_{11} > p_{22}$.

We are able to assess whether the microstructure assumptions underlying the proposed model are justified by looking at the estimates of m_0 and m_1. If m_1 is close to zero that would indicate that the information arrival does not affect "informed traders" and the model would be reduced to a simple stochastic volatility model (Jacquier et al., 1994). From the estimates in Table 1 we can conclude that the participants in this market do react to new information as, on average, 50% of trades are related to the arrival of news when $S_t = 2$ (high information state). When $S_t = 1$ (low information state) this proportion goes down to 15% with m_0 accounting for the remainder in both cases. These numbers are very similar to the findings in the finance literature as reported by Andersen (1996) and Watanabe (2000).

We would like to underscore an important point: the fact that most campaign events, particularly in mid-August to mid-September, had little or no new information is not saying that the campaign did not have any impact. Rather, it is saying that given the cast of characters and the political and economic context, the campaign itself played out predictably. It is true that during the month of August, there was an elevated state of information on just seven days, for a total of under 70 hours (or 22 periods with our measure), and July had just nine days with elevated information, for a total of just over 70 hours (or 24 periods with our measure). Despite a number of high-profile news stories during August, particularly the Swift Boat Veterans for Truth ads, the investors in the market felt that the underlying context of the election did not change very much. However, information that is not unexpected could, nonetheless, have been crucial to the electoral outcome.

Electoral outcomes can be thought to be composed of three pieces: context (economy, partisanship), standard campaign events, and election year-specific events. First, we will address the events specific to the 2004 presidential election. We will then present evidence regarding the effects of standard campaign events like the nominating conventions. A list of the most influential events appears in Table 2 and are marked by the red vertical lines in Figure 3.

4.1 National Security and the 2004 Election

Many people have described 2004 as a "national security" election. Opinion polls showed that people were highly concerned about the issue–in an ABC News[5] poll taken in early October, nearly half the respondents reported that either terrorism or the war in Iraq were the most important issues for the next president to deal with, and more importantly, made their vote

[4]To collect the data, we recorded the headlines on the front page of the site at each 3 hour interval. Major stories remained in the headlines for many periods. In theory, smaller stories could appear and disappear in-between our data collection points. However, if that is the case, we find it unlikely that that could be driving the markets.

[5]ABC News Poll. October 7, 2004. Data provided by The Roper Center for Public Opinion Research, University of Connecticut. [USABC.100804.R5]

choice while thinking about that issue (Newsweek, 10/29; Institute of Politics, 2006)[6]. Our analysis of the information flow during the campaign corroborates this: a number of national security-related issues significantly influenced the flow of information in the market. We will discuss several of these: the Swift Boat Veterans for Truth advertising campaign, the CBS story about Bush's National Guard service, the news about missing weapons in Iraq, and Osama bin Laden's taped message released just before the election.

Swift Boat Veterans for Truth (SBVT) Ads. After the 2004 election, a new verb appeared in political discussions: to "swift boat" someone meant to damage his candidacy with unfair and 'dirty' campaigning. The term itself may be useful, but is it accurate to apportion so much impact to the Swift Boat Veterans for Truth in the 2004 election? Our analysis suggests that while there was an increase in information flow at two points during the height of the story, the influence of the group was actually much less than is conventionally believed.

Though a book had been released months before detailing them, allegations made by the Swift Boat Veterans for Truth's that John Kerry had not earned his medals in the Vietnam War became a major news story on August 4, when SBVT first aired an advertisement. Because SBVT did not have much money to buy air time for their first ad, their strategy was to generate free media through the outlandish claims made in the ad. Needless to say, the strategy worked. Their first ad buy was for $500K, though after they gained notoriety, they significantly increased fundraising and by the end of the campaign, they had spent about $19M on television advertising (Institute of Politics, 2006).

News of the first advertisement regarding Kerry's medals elevated the information state for two days. Several days later, on August 9, the Drudge Report and others began questioning whether John Kerry was in Cambodia on Christmas 1969, as he previously claimed. It was later determined (and Kerry confirmed) that he was in Cambodia, but he was there in January 1970, not on Christmas. In retrospect, a rather unimportant point, but the markets reacted: the volatility increased on August 9 and 10. SBVT remained a major news story for most of August and has since made it into the vernacular, but the "information" from the SBVT allegations was apparently fully incorporated in approximately 36 hours (12 of our periods) over four days. There is another spike in the information flow on August 12 and 13, though there does not appear to be specific information about the SBVT allegations during that time. Even if that increase in volatility is attributed to SBVT, the information was incorporated into the market in just about 50 hours. It should be noted that the debates increased the information flow for a longer period of time than did the SBVT story (see below) and that the story regarding Bush's National Guard service (see below) increased the information flow for about the same amount of time. Additionally, though it is not the point of this paper, it is interesting to note that the likelihood that Bush would win the election did not change much from the beginning of August to the beginning of the Republican convention on August 30; the likelihood went from about 54% to about 56%. At the very least, because the change in the probability of Bush winning reelection was slight and our analysis suggests that the information was incorporated relatively quickly, it seems incorrect to describe Kerry's loss as a result of the "swift-boating" that took place in August 2004.

CBS Story on Bush's National Guard Story. John Kerry's military service was discussed frequently (by his campaign, his supporters, and also his detractors) throughout the primary and general election campaign. However, George W. Bush's military service had largely been ignored until September 8, when CBS News aired a story on "60 Minutes" that Bush had not properly reported for duty while he was in the National Guard. Dan Rather's reporting highlighted documents allegedly written by Bush's commander in the National Guard. The story developed over 12 days, with conservative bloggers immediately questioning the story and the documents and investigating the print on 1970s typewriters. It was several days into the story that the information state increased, on September 13, which coincided with the first reports by major news organizations that the papers were forged. The information flow remained high for most of the time between the 13^{th} and the 20^{th} and peaked on Sunday, September 19, when

[6]Princeton Survey Research International/Newsweek Poll. October 29, 2004. Data provided by The Roper Center for Public Opinion Research, University of Connecticut. [USPSRNEW.103004.R15]

the "60 Minutes" producer wrote an editorial in the *Washington Post* that stated the story was incorrect. It remained high until Dan Rather released a statement on the evening of September 20 admitting that the story was not properly reported and that the evidence was false. As this story developed, the volatility in the market rose and fell several times, indicating that new information was impacting investors and that it took some time for the potential consequences of the story to become clear.

Missing explosives in Iraq. Another national security issue burst into the campaign one week before Election Day. On Monday, October 25, the *New York Times* broke the story that 380 tons of powerful explosives vanished from an Iraq storage depot when the cache was supposed to be secured by American troops. Early that morning, when the Drudge Report initially reported it (presumably after it had appeared on the *New York Times* website), the information state increased immediately. Because this story went along with Senator Kerry's accusations that President Bush had incompetently handled the Iraq war after the initial invasion and the news was completely unexpected, it was viewed as informative. This heightened state of information lasted for three days.

Osama bin Laden videotape. The last days of the campaign showed an elevated level of volatility. On the Friday before the election (October 29), Al-Jazeera television showed a taped message from Osama bin Laden. After being authenticated by the FBI, American television stations began showing it. Bin Laden did not directly address the election, but he did say that, *"Your security is not in the hands of Kerry or Bush or al Qaeda. Your security is in your own hands. Any nation that does not attack us will not be attacked."* [7] When the ad was first played, the information flow increased dramatically; evidently people thought that bin Laden's discussion of 9-11 could influence the election by raising the salience of the issue. The information state remained high throughout the last three days of the campaign, in part because of the tape but presumably also because of new polls and reports of frenzied campaigning.

4.2 Effects of Standard Campaign Events

Scholars have found that nominating conventions (Campbell et al., 1992; Holbrook, 1996) and presidential debates (Lanoue, 1992; Benoit et al., 2003; Hillygus and Jackman, 2003) regularly impact voters. This is in part due to the fact that these events are highly publicized. Presidential debates are watched by millions of people; in 2004, 62.5 million viewers tuned into the first debate, 46.7 million watched the second, and 51 million people watched the third[8]. Viewership of nominating conventions has declined, but the candidate's acceptance speeches are still viewed by millions of people. It is debated whether or not the choice of a vice presidential candidate has an influence on the outcome, but at the least we can agree that the selection is a high-volume news event. Our measure of information flow reveals that the John Edwards' selection as John Kerry's running mate was viewed as important by investors and that the debates contained a significant amount of information. For the most part and contrary to prior findings, we show that the nominating conventions did not alter expectations about the election's outcome.

Vice Presidential Nomination. Political observers begin discussing possible vice presidential candidates early in the nominating process. Though it is unclear whether or not the vice presidential candidate adds much to the party's ticket, the conventional wisdom suggests that it can potentially impact election outcomes by winning additional states or bringing the party together after a nomination battle (Mayer, 2000). Our estimates of information flow suggest John Kerry's announcement that he would pick John Edwards, former senator from North Carolina, was considered to have a potential impact on the election's outcome.

[7] "Bin Laden: *'Your security is in your own hands'*". CNN, Friday, October 29, 2004 Posted: 10:05 PM EDT. `www.cnn.com/2004/WORLD/meast/10/29/bin.laden.transcript/`, viewed February 17, 2008.

[8] This is in comparison the 15.2 million who watched Fox's coverage of baseball playoffs. Though we find this heartening for democracy, it is nonetheless shocking. We have no idea why anyone would choose a debate over a Red Sox-Yankees game.

In 2004, after a strong second-place showing in the Iowa caucus, John Edwards was viewed as affable and a charismatic speaker. More importantly, he could perhaps help Kerry win the support of southern voters. On July 6, the day that the Kerry campaign announced that Edwards would be the vice presidential candidate, the information state was heightened. Investors thought that this could affect the outcome of the election. Interestingly, Edwards' selection was incorporated into expectations relatively quickly; the information state remained elevated for only 18 hours.

Nominating Conventions. Presidential nominating conventions have become highly choreographed media spectacles, though many people do watch parts of the four day conventions. Prior research suggests that party conventions do move voters (e.g. Holbrook, 1996). These findings are based on polls taken just before and just after conventions. Because our information flow variable is continuous and responds immediately to new information, we are able to better investigate the dynamics associated with conventions.

In 2004, the Democratic convention was held in Boston, Mass. on July 26-29 and the Republican convention was held in New York City on August 30-September 2. Our findings suggest that only one small part of one convention affected expectations about the ultimate election outcome: Kerry's acceptance speech on July 29 started an information incorporation period. The information state remained elevated for a full day[9]. At first, the finding that conventions had little effect seems somewhat surprising, though after further consideration, it makes sense: expected information has previously been incorporated into the market's prices. After four years, there was little that voters could learn about Bush or his policies and nearly all had opinions of him. The fact that volatility did not increase in demonstrates that the Republican convention did not contain any unexpected information.

In contrast, John Kerry was relatively unknown before his presidential bid; before the Democratic convention many people did not know much about Kerry, his background, or his policies. His speech was meant as an introduction to the candidate, and evidently investors thought that that introduction was, in fact, informative. Before the convention, Kerry's chances to win the popular vote were about even, and they remained around 50% after his speech. An analysis of simply the mean price, or probability, of a futures contract would suggest that the convention was completely irrelevant. This volatility measure suggests that despite the fact that there was not a significant change in price, people did learn about Kerry from the convention. The information acquired in this period likely conditioned the effect of later information and was thus indirectly important.

Both campaigns note that a key moment during the campaign occurred during the Republican convention (Institute of Politics, 2006). On September 1, Chechen terrorists invaded an elementary school in the Russian town of Beslan; after the Russian government stormed the school several days later, 200 people, mostly children, were killed. The campaigns argued that this reinforced Bush's message that the world is dangerous. Our analysis finds no such impact. Clearly, this incident could not possibly have been forecast by investors in the market, so what explains this? We can only speculate that because people already placed the issue of terrorism as a high priority for the next president, a certain ceiling effect existed; if the issue was ranked as the most important, it would be hard for the salience to increase very much.

Debates. There is a significant amount of evidence suggesting that televised presidential debates impact elections (Geer, 1988; Benoit et al., 2003), though *"their impact has fluctuated, from inconsequential to decisive"* (Graber, 1993). Even if debates do not change the winner of the election, citizens can and do learn from them (Lanoue, 1992; Holbrook, 1994; Miller and MacKuen, 1979). Our research confirms that debates provide a significant amount of information. In 2004, there were three presidential debates: on Thursday, September 30; Friday, October 8; and Wednesday, October 13. Bush was widely viewed as performing poorly in the first debate (e.g., *Los Angeles Times*, 10/2), though the reviews of the second and third debates were mixed. Each debate increased the volatility of the market. In particular, the first and

[9]We assume that this elevated information flow was due to Kerry's speech, and not to the failure of the balloons to drop immediately after his speech.

third debates led to an elevated information state for several days after the event. This extended length of volatility fits with findings that post-debate news analysis conditions people's responses to debates (Hwang et al., 2007).

The volatility surrounding the days of the debates highlights the difficulty isolating the cause of that increased information. The information flow increased immediately following the first debate, so it seems reasonable to conclude that that increase was caused by the debate. However, while the debate was the most notable event that occurred from September 30 to October 4 (when the volatility died down), a number of other plausible stories can be told for why the information flow was high three days later, on October 3^{rd}.

Similarly, the level of volatility increased more than a day before the third debate and continued for more than three days after. What information was the market incorporating? In that debate, when asked whether he believed homosexuality was a choice, Kerry replied, *"Dick Cheney's daughter, who is a lesbian, she would tell you that she's being who she was."* Both liberals and conservatives found the comment awkward and odd and thought he should not have said it. Was it that comment, or the fact that the polls moved only slightly over the course of the three debates, that caused the increased volatility? Or was it some other event? Days after that debate, the popular summary source of political information, ABC's *The Note*, commented:

> *"Because the polls moved slightly in Bush's favor after the debates, there must be a reason, and the only two reasons [the media] have been able to come up with are (i) the Mary Cheney remark; and/or (ii) the nation, having considered the totality of the debate round robin, decided it wanted a steady, likeable leader – rather than a voluble debating champion – to be the honcho of the free world. We have no idea if those are the reasons, but that's the best anyone has offered that has smacked against our ears."*

The problem of pinpointing causality in this kind of case seems, to us, intractable. Our method points to the fact that the debates added information to campaign, but we cannot apportion how much is due to comments during the debate, media analysis immediately following the debate, or even to separate events that happen the next day. In July and August, when less is going on in the campaign, it is easier to isolate events. By mid-October, so many polls are being released, so much money is being spent, and so many people are paying attention that it is much more difficult to make any kind of causal statement.

5 Concluding Remarks

The approach we develop here provides a novel and useful way to investigate the effects of campaigns. The information flow measure allows us to recognize key moments and separate the anticipated, normal dynamics of campaigns from the unexpected occurrences. By using data from prediction markets and a model that attempts to characterize the reactions of market participants to new information, we are able to address an as-yet unresolved question in political science. We build on a vast literature in finance and economics and propose a model with two volatility regimes that simplifies the task of associating events with periods of high information. The development and estimation of these models is greatly facilitated by the use of the Bayesian framework.

We identify a number of events that increased the information flow in the 2004 campaign. Among the standard campaign events, we find that the selection of the vice presidential candidate and the televised debates increased volatility in the prediction market. John Kerrys acceptance speech at the Democratic nominating convention also increased volatility though neither the rest of the convention or any of the Republican convention introduced new information. A number of events specific to the 2004 election were also viewed as informative by investors. In particular, stories related to national security-related issues increased volatility. These include the report that explosives vanished in Iraq under the U.S.'s watch, the CBS story about Bush's National Guard service and its subsequent retraction, and the release of the bin Laden tape a few days before the election. We find evidence that the Swift Boat Veterans for Truth ads that were critical of John Kerry's military service record were somewhat informative, but not nearly as central to the outcome of the 2004 campaign as is popularly believed.

We would like to reiterate two important points about this project. First, our goal was to

develop a quantitative way to objectively describe the information level over the course of a presidential campaign. The price of the Bush wins the popular vote futures contract chronicles the likelihood of Bush winning in 2004, and clearly whether or not certain events or certain classes of events impact that likelihood is an important question. However, the question of whether an event is informative is related to, but distinct from, the question of what effect a given event has on a particular election. That is, certain events can contain significant information even if they do not alter the outcome in a particular year. For example, if there is an overriding factor in the electionlike a bad economythen a candidate's acceptance speech at the party convention may not change the likelihood that the incumbent party will lose. Nonetheless, speeches at conventions do contain information, and in a year where the economy was not so salient, that speech could alter the election outcome. Put differently, some information that comes out during the campaign may be trumped by other information or considerations and thus not impact the vote percentage of a candidate. But that event does, in fact, have information in it. In a year with a different underlying political context, that same event may move vote totals significantly. Because our model focuses on the information level itself, we hope our model can be used to develop a more a general theory of campaign event effects over time.

Second, when we find relatively low levels of information in some given period of time, we are not saying that the campaign itself during those periods is irrelevant. We certainly believe that campaigns matter. The expectations about what will happen on election day are dependent on our prior understanding of how campaigns generally progress. When we find low levels of information, we are only arguing that those are periods where people are not learning anything new about the election.

We believe our model is a promising technique for research on campaigns and elections. We will be able to test it and our theory more explicitly in the future because prediction markets grow in popularity in each election cycle. Intrade has had contracts on the party nominees and the winner of the 2008 general election trading with high volume for months. Going forward with this research, we would like to find ways to make stronger causal claims. The question of causality looms large in the social sciences, and this work is no exception. We want to extend our modeling approach in an attempt to more clearly isolate the effects of different events. Addressing that becomes both more difficult and more important in the last weeks of the campaign. The 2008 Democratic primary season provides a particularly useful testing ground for us; because partisanship cannot be the deciding factor between two candidates in a primary, we are likely to see more dramatic movement in the markets during the campaign as the situation progresses.

There certainly is no shortage of political pundits, campaign strategists, and journalists willing to declare "what the election was about" and highlight the "turning point" in the campaign, the event or story that pushed the campaign towards its ultimate result. Undoubtedly, pundits and politicos are sometimes correct; but often they are not. When we explain election outcomes with a compact anecdote, we can both overlook events that were actually important and overstate our confidence in our conclusions. A Bayesian modeling approach to this problem, combined with theoretical insight from other fields, gives us greater leverage on a question that is of both academic interest and practical significance–it helps us to understand the dynamics and responsiveness of our democratic institutions.

Appendix A: Model and Computation

We complete the model described in (1) with the following priors: $\alpha_i \sim N(m_{\alpha_i}, C_{\alpha_i})$, $\beta \sim N(m_\beta, C_\beta)\mathbf{I}_{\{-1<\beta<1\}}$, $\tau^2 \sim IG(a_\tau, b_\tau)$, $p_{ii} \sim Be(a_i, b_i)$, $m_0 \sim Ga(a_{m_0}, b_{m_0})$ and $m_1 \sim Ga(a_{m_1}, b_{m_1})$. Let $\theta = \big(m_0, m_1, \alpha_1, \alpha_2, \beta, \tau^2, P\big)$. Based on the following decomposition of the joint posterior distribution

$$p\left(I_{1:T}, S_{1:T}, \theta | R_{1:T}, V_{1:T}\right) \propto p\left(R_{1:T}|I_{1:T}\right) p\left(V_{1:T}|I_{1:T}, \theta\right) p\left(I_{1:T}|S_{1:T}, \theta\right) p\left(S_{1:T}|\theta\right) p\left(\theta\right) \qquad (5)$$

we can generate posterior samples thought a MCMC schemes defined by the full conditionals described below.

5.1 Drawing θ

Full conditionals for θ can be easily obtained as follows:

$$(\alpha_i|\cdot) \sim N\left\{\left(\frac{T_i}{\tau^2}+\frac{1}{C_{\alpha_i}}\right)\left(\frac{\sum_{t:S_t=i}\lambda_t-\beta\lambda_{t-1}}{\tau^2}+\frac{m_{\alpha_i}}{C_{\alpha_i}}\right),\left(\frac{T_i}{\tau^2}+\frac{1}{C_{\alpha_i}}\right)^{-1}\right\}; \tag{6}$$

$$(\beta|\cdot) \sim N\left\{\left(\frac{\sum_{t=1}^{T}\lambda_{t-1}^2}{\tau^2}+\frac{1}{C_\beta}\right)\left(\frac{\sum_{t=1}^{T}(\lambda_t-\alpha_{S_t})\lambda_{t-1}}{\tau^2}+\frac{m_\beta}{C_\beta}\right),\left(\frac{\sum_{t=1}^{T}\lambda_{t-1}^2}{\tau^2}+\frac{1}{C_\beta}\right)^{-1}\right\}; \tag{7}$$

$$(\tau^2|\cdot) \sim IG\left(a_\tau+\frac{T}{2} \quad , \quad b_\tau+\frac{1}{2}\sum_{t=1}^{T}\lambda_t-\alpha_{S_t}-\beta\lambda_{t-1}\right); \tag{8}$$

$$(p_{11}|\cdot) \sim Be\left(a_1+n_{11}, b_1+n_{12}\right); \tag{9}$$

$$(p_{22}|\cdot) \sim Be\left(a_2+n_{22}, b_2+n_{21}\right); \tag{10}$$

$$p(m_0|\cdot) \propto \exp\left[-\left(b_{m_0}+T\right)m_0+\sum_{t=1}^{T} log\left\{m_0^{a_{m_0}-1}\left(m_0+m_1I_t\right)^{V_t}\right\}\right]; \tag{11}$$

$$p(m_1|\cdot) \propto \exp\left[-\left(b_{m_1}+\sum_{t=1}^{T}I_t\right)m_1+\sum_{t=1}^{T} log\left\{m_1^{a_{m_1}-1}\left(m_0+m_1I_t\right)^{V_t}\right\}\right], \tag{12}$$

where $\lambda_t = log(I_t)$, T_i represents the number of observations allocated to the i^{th} state and n_{ij} are counts of transitions from state i to j. Sampling from the above distributions is straightforward. The only minor complication is the Metropolis step necessary for sampling m_0 and m_1. We use a random-walk metropolis with normal proposals. Initial values where defined based on the MLE of both parameters in a Poisson regression where $I_{1:T}$ was fixed at the posterior mean from a standard two stage Markov-switching stochastic volatility model where an efficient Gibbs sampler can be implemented. These estimates of the information flow are also used as initial values for $I_{1:T}$.

5.2 Drawing the States $I_{1:T}$ and $S_{1:T}$

Due to its Markovian property and the decomposition in (5) the full conditional for $S_{1:T}$ can be written as:

$$\begin{aligned} p\left(S_{1:t}|\cdot\right) &= p\left(S_T|\lambda_{1:T},\theta\right)\prod_{t=T-1}^{1}p\left(S_t|S_{t+1},\lambda_{1:T},\theta\right) \\ &\propto p\left(S_T|\lambda_{1:T},\theta\right)\prod_{t=T-1}^{1}p\left(S_{t+1}|S_t\right)p\left(S_t,\lambda_t,\theta\right). \end{aligned} \tag{13}$$

This allows for a block sampling of $S_{1:T}$ through a *forward filtering backward sampling* algorithm as described in Kim and Nelson (1999). This is a simple task as S_t is a discrete random variable and obtaining the distributions defining (13) is always possible.

Sampling of I_t is the only "hard" step in the MCMC. The full conditional $p(I_t|I_{-t},\cdot)$ takes the following form

$$p\left(I_t|I_{-t},\cdot\right) \propto \exp\left(-\frac{1}{2I_t}R_t^2\right)\exp\left(m_0+m_1I_t\right)\left(m_0+m_1I_t\right)^{V_t}\times \tag{14}$$

$$\times\exp\left\{-\frac{1}{2\tau^2}\left(\lambda_t-\alpha_{S_t}-\beta\lambda_{t-1}\right)^2\right\}\exp\left\{-\frac{1}{2\tau^2}\left(\lambda_{t+1}-\alpha_{S_{t+1}}-\beta\lambda_t\right)^2\right\} \tag{15}$$

Joint samples from the full conditional of $I_{1:T}$ cannot be generated without approximations such as de Jong and Shephard (1995). In this work we implement a *single-move* MCMC and generate draws from $p(I_t|I_{-t}, \cdot)$ as in Carlin et al. (1992). Once again, random-walk metropolis was used. It is our experience that this simple approach performed well enough and our MCMC presented satisfactory convergence diagnostics. Alternative and more efficient methods are described in Abanto-Valle et al. (2008).

References

Abanto-Valle, C., H. Migon, and H. Lopes (2008). Bayesian modeling of financial returns: a relationship between volatility and trading volume. *Discussion paper, Federal University of Rio de Janeiro*.

Andersen, T. (1996). Return volatility and trading volume: An information flow interpretation of stochastic volatility. *Journal of Finance 51*, 169–204.

Benoit, W. L., G. Hansen, and R. Verser (2003). A meta-analysis of the effects of viewing us presidential debates. *Communication Monographs 70*, 335–350.

Berg, J., F. Nelson, and T. Rietz (2003). Accuracy and forecast standard error of prediction markets. *Discussion paper, Tippie College of Business, University of Iowa*.

Berg, J. and T. Rietz (2006). The iowa electronic market: Lessons learned and answers yearned. In R. Hahn and P. Tetlock (Eds.), *Information Markets: A New Way of Making Decisions*. AEI Press.

Bernhard, W. and D. Leblang (2006). *Democratic Processes and Financial Markets: Pricing Politics*. Cambridge University Press, New York.

Brown, G. and J. Hartzell (2001). Market reaction to public information: The atypical case of the boston celtics. *Journal of Financial Economics 60*, 333–370.

Campbell, J., L. Cherry, and K. Wink (1992). The convention bump. *American Politics Research 20*, 287–307.

Carlin, B., N. Polson, and D. Stoffer (1992). A Monte Carlo approach to nonnormal and nonlinear state-space modeling. *Journal of the American Statistical Association 87*, 493–500.

Carvalho, C. and H. Lopes (2007). Simulation-based sequential analysis of markov switching stochastic volatility models. *Computational Statistics and Data Analysis 51*, 4526–4542.

Chan, W. (2003). Stock price reaction to news and no-news: drift and reversal after headlines. *Journal of Financial Economics 70*, 223–260.

de Jong, P. and N. Shephard (1995). The simulation smoother for time series models. *Biometrika 82*, 339–350.

Dubinsky, A. and M. Johannes (2006). Earnings annoucement and equity options. *Working paper*.

Eraker, B., M. Johannes, and N. Polson (2003). The impact of jumps in volatility and returns. *The Journal of Finance 58*, 1269–1300.

Fleming, J., C. Kirby, and B. Ostdiek (2006). Stochastic volatility, trading volume, and the daily flow of information. *The Journal of Business 79*, 1551–1590.

Forsythe, R., M. Frank, V. Krishnamurthy, and T. Ross (1998). Markets as predictors of election outcomes: Campaign events and judgement bias in the 1993 ubc election stock market. *Canadian Public Policy 24*, 329–351.

Forsythe, R., F. Nelson, G. Neumann, and J. Wright (1992). Anatomy of an experimental political stock market. *American Economic Review 82*, 1142–1161.

Forsythe, R., T. Rietz, and T. Ross (1999). Wishes, expectations and actions: a survey on price formation in election stock markets. *Journal of Economic Behavior and Organization 39*, 83–110.

Gallant, A., P. Rossi, and G. Tauchen (1992). Stock prices and volume. *Review of Financial Studies 5*, 199–242.

Geer, J. G. (1988). The effects of presidential debates on the electorates preferences for candidates. *American Politics Quarterly 16*, 486–501.

Graber, D. (1993). *Mass Media and American Politics* (4th ed.). Congressional Quarterly Press.

Granberg, D. and E. Brent (1983). When prophecy bends: The preference-expectation link in us presidential elections, 1952-1980. *Journal of Personality and Social Psychology 45*, 477–491.

Gronke, P. and J. Brehm (2002). History, heterogeneity, and presidential approval: a modified arch approach. *Electoral Studies 21*, 425–452.

Grossman, S. and J. Stiglitz (1980). On the impossibility of informationally efficient markets. *American Economic Review 70*, 393–408.

Halperin, M. and J. F. Harris (2006). *The way to win: taking the White House in 2008*. New York: Random House.

Hartog, D. and N. Monroe (2008). The value of majority status: The effect of jeffords's switch on asset prices of republican and. *Legislative Studies Quarterly 33*, 63–84.

Hillygus, D. and S. Jackman (2003). Voter decision making in election 2000: Campaign effects, partisan activation, and the Clinton legacy. *American Journal of Political Science 47*, 583–596.

Holbrook, T. (1994). The behavioral consequences of vice-presidential debates: Does the undercard have any punch? *American Politics Research 22*, 469–482.

Holbrook, T. (1996). *Do Campaigns Matter?* Sage Publications Inc.

Hwang, H., M. Gotlieb, S. Nah, and D. McLeod (2007). Applying a cognitive-processing model to presidential debate effects: Postdebate news analysis and primed reflection. *Journal of Communication 57*, 40–59.

Jacquier, E., N. Polson, and P. Rossi (1994). Bayesian analysis of stochastic volatility models. *Journal of Business and Economic Statistics 12*, 371–389.

Jennings, R. and L. Starks (1985). Information content and the speed of stock price adjustment. *Journal of Accounting Research 23*, 336–350.

Kim, C. and C. Nelson (1999). State-space models with regime switching: Classical and gibbs-sampling approaches with applications.

Lamoureux, C. and W. Lastrapes (1990). Persistence in variance, structural change and the garch model. *Journal of Business and Economic Statistics 8*, 225–234.

Lanoue, D. (1992). One that made a difference: Cognitive consistency, political knowledge, and the 1980 presidential debate. *Public Opinion Quarterly 56*, 168–184.

Lopes, H. and C. Carvalho (2007). Factor stochastic volatility with time varying loadings and markov switching regimes. *Journal of Statistical Planning and Inference 137*, 3082–3091.

Mayer, W. G. (2000). *In pursuit of the White House 2000: how we choose our presidential nominees*. New York: Chatham House Publishers.

Miller, A. and M. MacKuen (1979). Learning about the candidates: The 1976 presidential debates. *Public Opinion Quarterly 43*, 326–346.

Oliven, K. and T. Rietz (2004). Suckers are born but markets are made: Individual rationality, arbitrage, and market efficiency on on an electronic futures market. *Management Science 50*, 336–351.

Patell, J. and M. Wolfson (1984). The intraday speed of adjustment of stock prices to earnings and dividend announcements. *Journal of Financial Economics 13*, 223–252.

Popkin, S. (1991). *The Reasoning Voter: Communication and Persuasion in Presidential Campaigns*. University of Chicago Press, Chicago.

Sauer, R. (1998). The economics of wagering markets. *Journal of Economic Literature 36*, 2021–2064.

So, M., K. Lam, and W. Li (1998). A stochastic volatility model with markov switching. *Journal of Business and Economic Statistics 16*, 244–253.

Uhlaner, C. and B. Grofman (1986). The race may be close but my horse is going to win: Wish fulfillment in the 1980 presidential election. *Political Behavior 8*, 101–129.

Watanabe, T. (2000). Bayesian analysis of dynamic bivariate mixture models: Can they explain the behavior of returns and trading volume? *Journal of Business and Economic Statistics 18*, 199–210.

Wolfers, J. and E. Zitzewitz (2004). Interpreting prediction market prices as probabilities. *Working Paper*.

Parameter	Posterior Mean	95% Interval
α_1	-0.369	(-0.395;-0.342)
α_2	-0.306	(-0.327;-0.284)
β	0.967	(0.927;0.983)
τ^2	0.019	(0.018;0.020)
m_0	3.936	(3.81;4.03)
m_1	28370	(25478;32005)
p_{11}	0.983	(0.932;0.991)
p_{22}	0.853	(0.801;0.912)

Table 1: Posterior Summary

July 6	Kerry picks Edwards as VP
July 29	Kerry's nomination speech
August 5	First Swift Boat Veterans for Truth advertisement airs
August 20	Second SBVT ad airs
September 8	Bush's National Guard service questioned in CBS story
September 19	CBS retracts the National Guard story and Dan Rather releases statement
September 30	First debate-Bush performs poorly
October 8	Second debate
October 13	Third debate-Kerry makes comment about Cheney's daughter
October 26	Story about missing explosives in Iraq in 2003

Table 2: Key Dates in 2004 Presidential Election

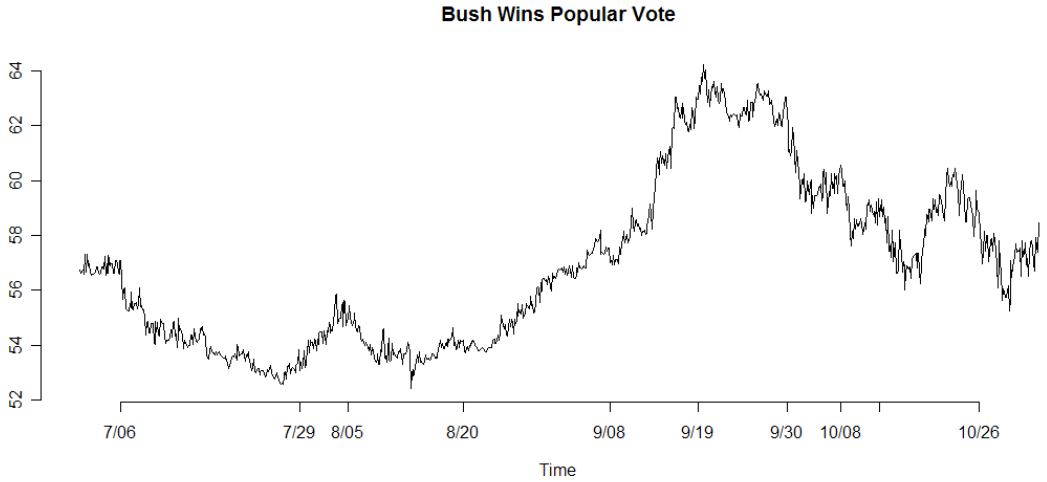

Figure 1: Closing price for 'Bush Wins the Popular Vote' at every 3 hour period. Source: Tradesports.com

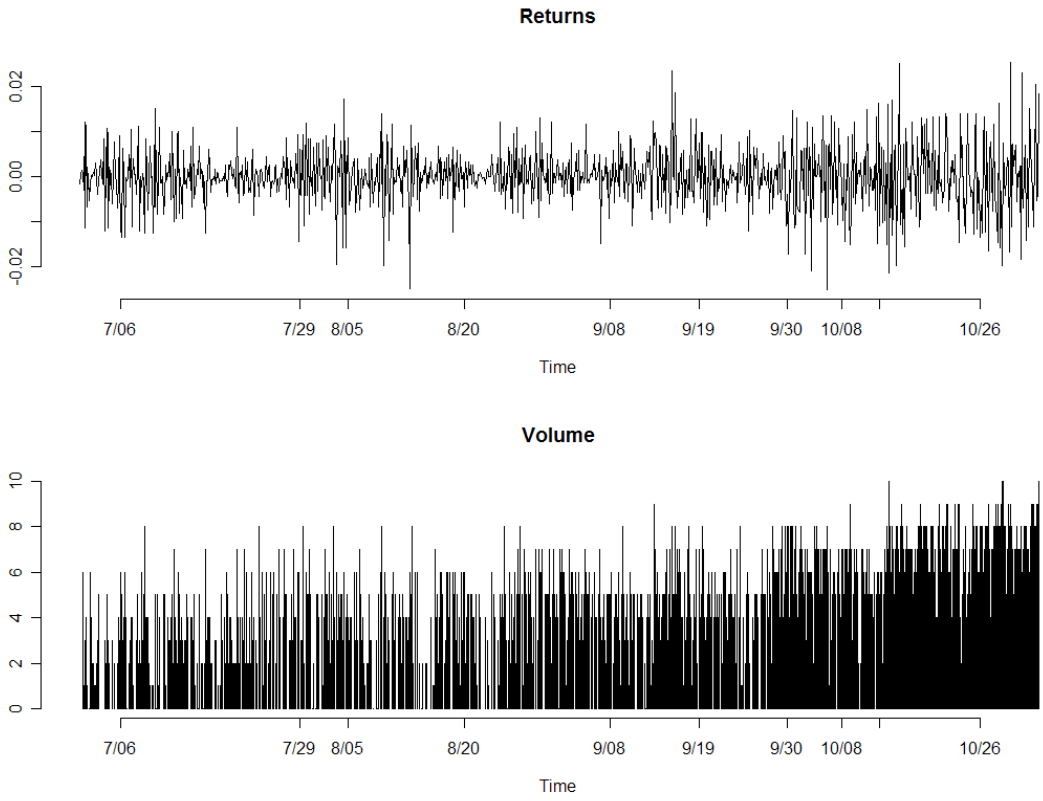

Figure 2: Returns and Trading Volume.

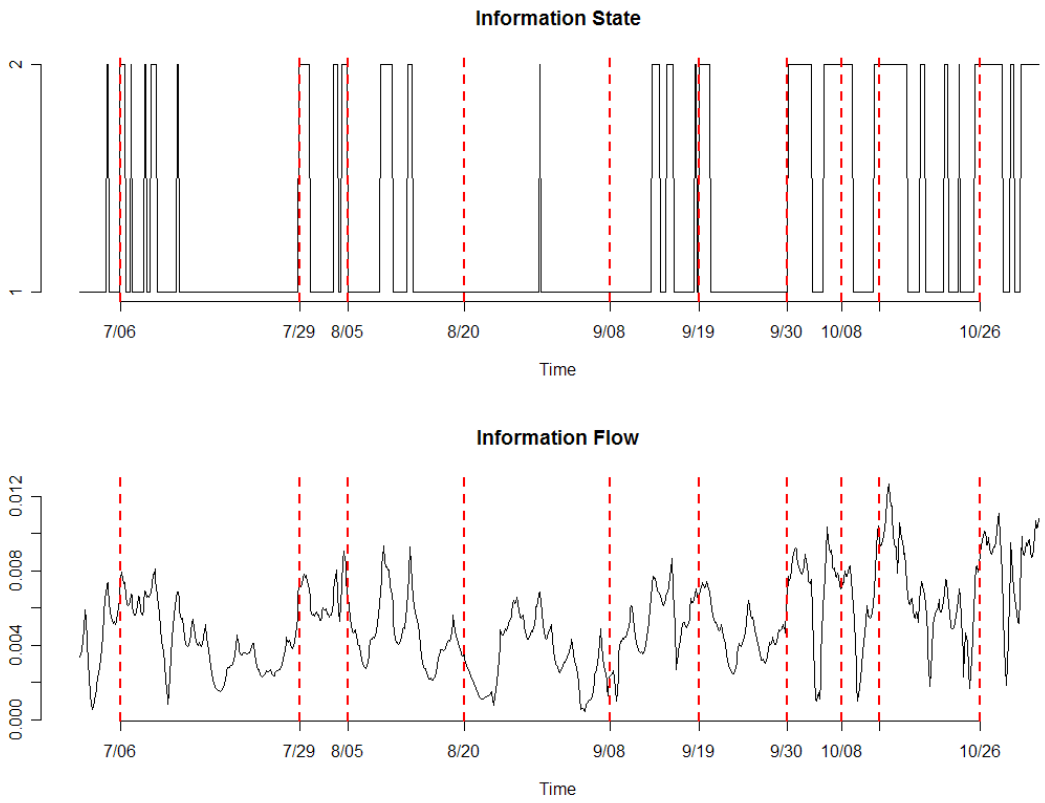

Figure 3: Posterior estimates of $S_{1:T}$ (top) and $\sqrt{I_{1:T}}$ (bottom). We define $\hat{S}_t = 2$ if the posterior $Pr(S_t = 2) > 0.5$. The red lines indicate the key events described in Table 2.