MANAGEMENT SCIENCE
Vol. 54, No. 8, August 2008, pp. 1381–1394
ISSN 0025-1909 | EISSN 1526-5501 | 08 | 5408 | 1381

DOI 10.1287/mnsc.1080.0872


# Modeling a Presidential Prediction Market

M. Keith Chen
Yale School of Management and the Cowles Foundation, Yale University, New Haven, Connecticut 06520, keith.chen@yale.edu

Jonathan E. Ingersoll, Jr.
Yale School of Management, Yale University, New Haven, Connecticut 06520, jonathan.ingersoll@yale.edu

Edward H. Kaplan
Yale School of Management, Yale School of Medicine, Yale School of Engineering and Applied Science, Yale University, New Haven, Connecticut 06520, edward.kaplan@yale.edu

Prediction markets now cover many important political events. The 2004 presidential election featured an active online prediction market at Intrade.com, where securities addressing many different election-related outcomes were traded. Using the 2004 data from this market, we examined three alternative models for these security prices, with special focus on the electoral college rules that govern U.S. presidential elections to see which models are more (or less) consistent with the data. The data reveal dependencies in the evolution of the security prices across states over time. We show that a simple diffusion model provides a good description of the overall probability distribution of electoral college votes, and an even simpler ranking model provides excellent predictions of the probability of winning the presidency. Ignoring dependencies in the evolution of security prices across states leads to considerable underestimation of the variance of the number of electoral college votes received by a candidate, which in turn leads to overconfidence in predicting whether that candidate will win the election. Overall, the security prices in the Intrade presidential election prediction market appear jointly consistent with probability models that satisfy the rules of the electoral college.

*Key words*: prediction market; stochastic model applications; U.S. presidential election; electoral college
*History*: Accepted by Linda V. Green, public sector applications; received May 2, 2007. This paper was with the authors 3 months for 2 revisions. Published online in *Articles in Advance* June 10, 2008.

## 1. Introduction

The outcome of the U.S. presidential election is of widespread interest and importance, yet there is considerable disagreement among experts and pundits when prospectively predicting election outcomes. Many forecasters focus on national polling data (Fair 2002, Wlezien and Erikson 2007), yet recent presidential elections have demonstrated the difficulty of translating national vote shares into electoral college outcomes (Erikson and Sigman 2000, Jackman and Rivers 2001, Kaplan and Barnett 2003). Even if state-by-state polling data are available, how should they be aggregated into an electoral college prediction?

With the 2004 presidential election, a new data source was introduced. An active Internet prediction market hosted at Intrade.com traded numerous securities anchored in different aspects of the election. The security prices generated by this market enable us to explore two things. First, by examining the internal consistency of security prices whose contracted underlying events are related by the electoral college rules governing presidential elections, we are able to see whether the Intrade prediction market prices can be interpreted as probabilities (Manski 2006, Wolfers and Zitzewitz 2007). Second, after demonstrating that such a probabilistic interpretation is reasonable, we evaluate alternative models for the probability distribution of the number of electoral college votes implied by the security prices. Such models produce an estimate of the probability of winning the presidency as a byproduct. We consider three models for the electoral college distribution: independence (which assumes that state-by-state election outcomes are independent, as in Kaplan and Barnett 2003); Fair's ranking model, which presumes maximal dependence across states (Fair 2007); and a new one-parameter stochastic diffusion model that nests both of these extremes while explicitly modeling the evolution of state-by-state probabilities.

This paper proceeds as follows. In §2, we briefly describe prediction market data, securities associated with the U.S. presidential election that are traded, and the issues involved with interpreting prices in these markets as probabilities. Section 3 presents a general model of the electoral college that highlights some of the relationships any specific electoral college model must satisfy, and §4 presents simple plausibility checks on the internal consistency of the Intrade prediction market data. In §5 we present two modeling benchmark extremes—Kaplan and Barnett's (2003) independence model (§5.1) and Fair's (2007) perfect correlation "ranking" model (§5.2)—before introducing our new diffusion model in §5.3. In describing these models, we present many of their formal properties (some for the first time in the literature), implications, and drawbacks.

Section 6 considers these models in light of the Intrade data available during the 2004 U.S. presidential election. We provide direct evidence that in this market, market-anticipated electoral outcomes across several states were highly correlated (§6.1). We then fit the diffusion model to the Intrade data in §6.2. In §7, we compare the three models with respect to their ability to reproduce the variability in the electoral college distribution evident in the data as well as the market-implied probability of winning the presidency; §8 concludes.

## 2. The Intrade Political Betting Market

In the run up to the 2004 presidential election, several website and online betting markets offered tradeable securities on the election's outcome. The largest of these betting markets was the Intrade (or Tradesports) betting market, located at http://www.intrade.com. Intrade allows users to both buy and sell contracts in units of $10. Buying any of these contracts pays $10 if the stated event occurs and $0 otherwise. For example, purchasing the contract, "George W. Bush to win the electoral votes of Florida" for a price of $6 yields a payoff of $10 if Bush wins Florida (for a profit of $4) and yields a loss of $6 if Bush loses Florida. The market runs continuously, so at any time after a contract is purchased it can be resold at the current price, unwinding the initial bet. All the security prices employed in this study are bid-ask midpoints observed daily at midnight U.S. eastern time.

The contracts that were offered included the following: Bush to win each of the 50 states or Washington, D.C.; Bush to win several different sets of states (for example, Bush to win California, Oregon, and Washington); Bush to win reelection; and several contracts of the form "Bush to win $\geq x$ electoral college votes." Many prediction contracts had high volume, with more than $1 billion (over the entire Intrade market) changing hands in the three years prior to the election; several of the more popular contracts contained tens of thousands of dollars of open interest at any given time.

The use of prices to deduce probability distributions in finance has a large literature. Latane and Rendleman (1976) were the first to propose using an option price to compute *implied volatility*, that is, the standard deviation of price returns. Bodurtha and Shen (1999) and Campa and Chang (1998), among others, have used similar methods to deduce the implied correlations between two underlying variables using individual derivative contracts written on each underlying variable as well as contracts written on both variables. Skintzi and Refenes (2005) discuss many aspects of calculating implied correlations. Banz and Miller (1978) and Breeden and Litzenberger (1978) demonstrated that a set of option prices could be used to infer the entire probability distribution. Although the interpretation of prediction market trading prices as market probability estimates is obvious,[1] their dependability has been debated (Manski 2006, Wolfers and Zitzewitz 2007). For an excellent overview of prediction markets as a whole and their structure, see Wolfers and Zitzewitz (2004).

Following Wolfers and Zitzewitz (2007), we assume that the trading price of a contract is an unbiased market estimate of the underlying event's likelihood, and we explore the reasonability of this assumption in §4. In the remainder of this paper we develop various models of both how the prices on these contracts should relate to each other and how they can be used to predict several features of the overall presidential election and its outcome.

## 3. Modeling the Electoral College

There are 51 "states" in the election, including Washington, D.C. (which we henceforth include as a state for expositional simplicity). We assume that all of the electoral college votes in a given state are allocated to the candidate who wins the most votes in that state. This is true by law in all states except Maine and Nebraska, and in practice neither state has ever split its electoral votes. We focus on a particular candidate and let $p_i$ denote the probability that this candidate will win the popular vote in state $i$. We take the market price on a contract for the candidate in question to win state $i$ as our estimate of $p_i$; thus, what follows can be thought of as an assessment given the data on a given day.

There are $v_i$ electoral college votes at stake in state $i$. We let the random variable $V_i$ denote that actual number of electoral college votes won by the candidate in

[1] If traders as a whole are indifferent to the risk inherent in the event described by the ticket, then the ticket price should be the probability of the event discounted at the risk-free rate until the time of the payout. The period of study includes only the last five months before the election, and T-bill rates in 2004 were about 1.2%, so the price would underestimate the probability by a factor of $(1.012)^{-5/12} - 1 \approx 0.5\%$. For example, an observed price of 0.50 would indicate a probability of 50.2%. If traders are not indifferent to the risk, then the discount rate would include a (positive or negative) risk premium as well.

state $i$, and via our "all-or-nothing" assumption note that the probability distribution of $V_i$ is given by

$$\Pr\{V_i = v\} = \begin{cases} p_i & \text{if } v = v_i, \\ 1 - p_i & \text{if } v = 0, \\ 0 & \text{for all other values of } v, \end{cases} \qquad i = 1, 2, \ldots, 51. \quad (1)$$

Let $T$ denote the total number of electoral college votes earned by our candidate in the election. Clearly

$$T = \sum_{i=1}^{51} V_i, \quad (2)$$

from which it follows that the expected number of electoral college votes earned by the candidate equals

$$E(T) = \sum_{i=1}^{51} E(V_i) = \sum_{i=1}^{51} p_i v_i. \quad (3)$$

Now define $p_{ij} = \Pr\{\text{Candidate wins } both \text{ states } i$ and $j\}$ in the election. The variance of the number of electoral college votes garnered by our candidate is then given by

$$\begin{aligned} \mathrm{Var}(T) = \sigma_T^2 &= \sum_{i=1}^{51} \mathrm{Var}(V_i) + 2\sum_{i<j} \mathrm{Cov}(V_i, V_j) \\ &= \sum_{i=1}^{51} p_i(1-p_i)v_i^2 + 2\sum_{i<j}(p_{ij} - p_i p_j)v_i v_j. \quad (4) \end{aligned}$$

By the rules of the electoral college, a candidate wins the presidency either by gaining a majority of the 538 total electoral college votes or—if no candidate receives an electoral college majority—by a state-by-state vote in the House of Representatives. Our empirical analysis addresses the 2004 election, where we assume that in the event of an electoral college tie, George Bush would have won the vote in the House.[2] Thus, we model the probability of Bush winning the presidency as

$$\Pr\{\text{Bush Wins Election}\} = \Pr\{T \geq 269\}. \quad (5)$$

[2] If no candidate receives a majority of the electoral votes, then the House of Representatives chooses the president from the top three recipients of electoral votes, with each state (not representative) casting one vote. We assumed that in 2004 only Bush and Kerry would receive electoral votes and that an electoral college tie of 269 votes each would have been resolved in Bush's favor in the House with 100% probability. This assumption is plausible because the 109th Congress, which was elected in 2004, was the one that would have broken the tie. In the House, Republicans controlled 31 states, Democrats controlled 15 states, and 4 states were equally divided. Furthermore, this breakdown would have been foreseen with high reliability because only seven states had any change in representative affiliations from the 108th Congress, and in only two states was there a change in majority—Illinois went from a 10–9 Republican majority to the opposite, and Texas went from an equal split to a 21–11 Republican majority.

Without further assumptions enabling the construction of the *probability distribution* of $T$, however, it is not possible to compute the win probability. We will explore alternative additional assumptions that do enable computing the distribution of $T$, but before doing so, we first consider some simple plausibility checks to see if the main data from the prediction market are consistent with the basic workings of the electoral college.

## 4. Plausibility Checks on Prediction Market Consistency

Figure 1 plots the expected number of electoral votes for George Bush over the days prior to the election as computed from Equation (3) by interpreting the daily state-by-state prices for the "Bush win" contracts in each state as market-derived probabilities of the underlying "Bush win" events. Also shown are the prices for "electoral college cutoff" contracts on Bush winning at least 250, 269 (i.e., winning the election) and 300 electoral college votes. Note first that the prices on the electoral college contracts are always ranked in the correct direction according to probability theory, that is, $\Pr\{T \geq 250\} \geq \Pr\{T \geq 269\} \geq \Pr\{T \geq 300\}$. Second, the prices on these contracts clearly move together, as would be expected from shifts in the probability distribution of $T$ as events unfold over time. Third, over most of the time period covered, on those days when $E(T) \geq 269$ (the horizontal line on the chart), the probability that Bush wins the presidency exceeds $1/2$ (also the horizontal line on the chart). And fourth, small swings in $E(T)$ in the vicinity of 269 produce much greater swings in $\Pr\{T \geq 269\}$, the probability of Bush winning the election, again consistent with probability theory.

As a second plausibility check, suppose that the distribution of $T$ is approximately symmetric and differentiable about its mean $\mu = E(T)$.[3] Then for various electoral college vote totals $x$ close to the mean value $\mu$, the probability of observing at least $x$ electoral college votes (equivalently the price on a contract that the candidate would receive at least $x$ electoral college votes) should be (by a Taylor expansion)

$$\Pr\{T \geq x\} \approx \int_x^\infty f_T(u)\,du \approx \tfrac{1}{2} + f_T(\mu)(\mu - x), \quad (6)$$

[3] Although this is not an unreasonable assumption given that the total number of electoral college votes won is the sum of 51 random variables, we recognize that there are circumstances where this assumption is unquestionably false, as when, for example, the chance that a candidate will win California's 55 electoral college votes is bounded well away from zero or one; see Kaplan and Barnett (2003) for an example of the resulting bimodal distribution.

**Figure 1** **Prediction Market Summary**

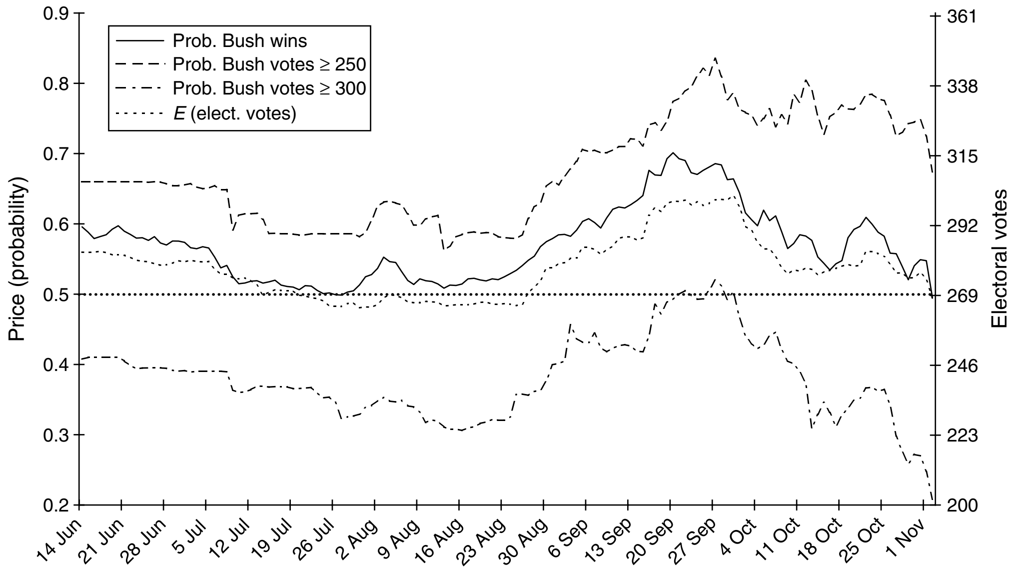


where $f_T(\cdot)$ is the approximating symmetric density of $T$. This suggests that if one plots the prices on electoral college cutoff securities ($\Pr\{T \geq x\}$) against the expected number of electoral college votes ($\mu$) on each day for different cutoffs $x$ close to $\mu$, one should see parallel lines offset due to the different values of $x$. Figure 2 reports such a plot for the three cutoffs closest to the market's mean sentiment (i.e., at $x = 250$, 269, and 300 electoral votes). Consistent with Equation (6), the plot does indeed suggest parallel lines with appropriate offsets. This is impressive, considering that $\mu = E(T)$ was computed from the daily prices associated with the 51 state-by-state "Bush win" contracts, again consistent with the interpretation of trading prices as probabilities.

Note that if the variance of the total number of electoral college votes is driven by uncertainty regarding how traders believe voters will respond to information as it unfolds during the run up to the election, one would expect to see greater variability in $T$ at earlier points in time than at times closer to the election. This means that $f_T(\mu)$ should increase over time whether $\mu$ itself changes or not, which in turn implies that the parallel lines should be further apart after significantly informative events. Figure 2 seems to confirm this intuition. We have separately plotted points from before and after the presidential debates, as these seem the most informative events during the campaign. Note that after the debates, $\Pr\{T \geq 250\}$ is much higher than it was before the debates, and

**Figure 2** **Electoral College Prices, Before and After Debates**

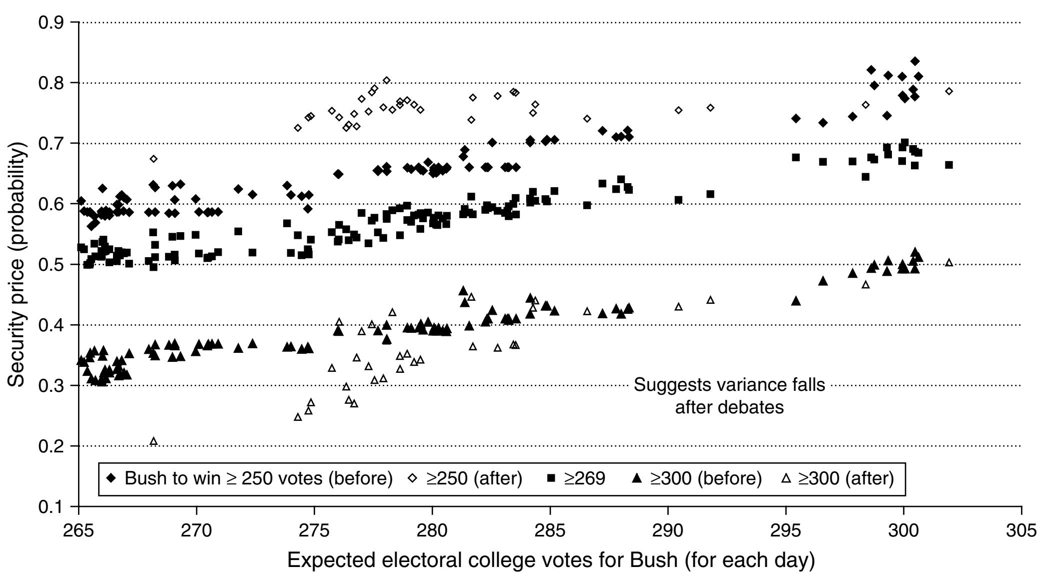

$\Pr\{T \geq 300\}$ is much lower after the debates than before. Both of these results are consistent with a decline in the variance of the number of electoral college votes after then debates: as the variance falls, more of the probability concentrates in the vicinity of $\mu$, which lowers the chance that $T$ exceeds values greater than $\mu$ while raising the chance that $T$ exceeds values smaller than $\mu$ (and note that after the debates, the daily value of $\mu$ typically fell between 275 and 285).

We do not believe that the plausibility checks above completely establish the consistency of the prediction market. However, had the market prices *failed* to satisfy these simple tests, there would have been little point to continuing our investigation.

## 5. Alternative Models of the Electoral College Distribution

As stated earlier, to compute the probability of winning the presidency from the state-by-state presidential win contracts requires additional assumptions to construct the probability distribution of $T$. The idea is to see which (if any) assumptions are able to replicate the distribution of observed prices of electoral victory over the days and months prior to the election. Success in this endeavor, we hope, will reveal an "as if" principle by which the market is processing information to assess the probability of various election outcomes.

### 5.1. An Independence Model

The simplest model, previously studied in the context of public opinion polls by Kaplan and Barnett (2003), can be defined by imposing the independence restriction

$$p_{ij} = p_i p_j, \tag{7}$$

which reduces the variance of the number of electoral college votes to

$$\mathrm{Var}(T) = \sum_{i=1}^{51} p_i(1-p_i)v_i^2. \tag{8}$$

Under this model, the exact probability distribution of $T$ can be derived via recursion (Kaplan and Barnett 2003). Arbitrarily rank order the states from 1 to 51 and define $T_k$ as the total number of electoral college votes received by our candidate when considering states 1 through $k$. The recursion follows by noting that

$$T_{k+1} = V_{k+1} + T_k \quad \text{for } k = 1, 2, \ldots, 50. \tag{9}$$

Thus, for any state $k$ and any number of electoral college votes $x$, we have

$$\Pr\{T_{k+1} = x\} = (1-p_{k+1})\Pr\{T_k = x\} + p_{k+1}\Pr\{T_k = x - v_{k+1}\} \tag{10}$$

for $k = 1, 2, \ldots, 50$ and $x = 0, 1, \ldots, \sum_{i=1}^{k+1} v_i$. That is, there are only two ways our candidate could garner exactly $x$ electoral college votes from the first $k+1$ states considered: either the candidate already had exactly $x$ votes from the first $k$ states and lost in state $k+1$, or the candidate had exactly $x - v_{k+1}$ votes from the first $k$ states and then won an additional $v_{k+1}$ votes by winning state $k+1$. Iterating this recursion yields the exact probability distribution of electoral college votes under independence, enabling the calculation of the probability of winning the presidency (which is again just $\Pr\{T \geq 269\} = \sum_{x=269}^{538} \Pr\{T = x\}$).

This model would make intuitive sense if traders believe that voters have largely made up their minds by the day in question and would thus correspond to the common polling question "If the election was held today, for whom would you vote?" In this case, state-by-state win probability estimates are comparable to polls with uncorrelated sampling errors; the probability of winning a state would then be the conditional probability that a candidate achieves more than 50% support in that state given the poll results (Kaplan and Barnett 2003).

### 5.2. A Ranking Model

The ranking model proposed by Fair (2007) works as follows: first rank order the states from highest to lowest probability of voting for the candidate (i.e., in descending order of the $p_i$s). Then impose the rank restriction assumption, which states that if a candidate eventually wins in state $i$ (which occurs with probability $p_i$), then the candidate must also win in all states $j$ for which currently $p_j \geq p_i$. This assumption further implies that if the candidate eventually loses in state $i$, then the candidate also loses in all states $j$ where $p_j \leq p_i$. The idea behind this model is that good news raises the probability of winning in all states, and bad news lowers the probability of winning in all states. This happens in a monotonic fashion so that the overall ranking across states is preserved, even after the election results are realized.

Mathematically, the ranking model imposes the restriction

$$p_{ij} = \min(p_i, p_j) \tag{11}$$

for any two states $i$ and $j$. Note that because $p_{ij} \leq \min(p_i, p_j)$ for *any* possible model, Equation (4) shows that this model *maximizes* the variance of the electoral college distribution, and in this sense the ranking model creates the most diffuse distribution of $T$ that is consistent with any set of state probability estimates. Also note that the ranking model assumptions imply that the probability of winning any subset of states equals the minimum of the state-specific win probabilities over that subset (e.g., $p_{ijkl} = \min(p_i, p_j, p_k, p_l)$).

The rank restriction assumption leads to a remarkable reduction in the number of possible outcomes to the election. Given that there are 50 states plus Washington, D.C., each to be won or lost, there are $2^{51} \approx 2.25$ million-billion different possible outcomes to the election in general. The rank restriction reduces the number of possible outcomes to just 52: our candidate loses all 51 states, our candidate wins only his most highly ranked state, only the first two most highly ranked states, ..., all 50 states plus Washington, D.C.

Although Fair (2007) posited the rank restriction assumption, he did not develop the consequences of this for the probability distribution of the number of electoral college votes. We do so here. Under the ranking assumption, if our candidate wins the $i$th ranked state, then he must win all the more highly ranked states 1 through $i-1$ as well. Thus, the probability distribution of the total number of electoral college votes $T$ satisfies

$$p_i = \Pr\left\{T \geq \sum_{j=1}^{i} v_j\right\}, \tag{12}$$

which leads to the 52 mass point distribution for $T$

$$\Pr\left\{T = \sum_{j=1}^{i} v_j\right\} = p_i - p_{i+1} \tag{13}$$

for $i = 0, 1, \ldots, 51$ with $p_0 \equiv 1$ and $p_{52} \equiv 0$.

Finally, the probability of winning the election under this model is given by finding that state $i^*$ that satisfies

$$\Pr\{T \geq 269\} = \max_i \left\{p_i \;\middle|\; \sum_{j=1}^{i} v_j \geq 269\right\}. \tag{14}$$

Fair (2007) noted that under the rank restriction assumption, the probability of winning the election reduces to the probability of winning a single *pivotal* state $i^*$. Algorithmically, to find the pivotal state, one orders the states from highest to lowest state-winning probabilities, scans down the list while tallying the cumulative electoral college votes, and notes the first state on the list for which the cumulative electoral college votes equals or exceeds 269. This state is the pivotal state $i^*$, and the probability of winning the election is just the probability $p_{i^*}$ of winning the pivotal state.[4]

[4] As first posited, Fair's (2007) ranking model is about election uncertainty on the morning of the election. There is no assumption that a single pivotal state remains the same throughout the campaign. However, when Fair examines the ranking assumption using Intrade data, he does extend his assumption in this way by assuming that the states are correctly ranked by *current* prices. In particular, the prices in Fair's (2007) Table 4 for the various "region-sweeping" contracts are computed as the minimum of the prices of the included states, and the price of the "election-win" contract is set equal to the price of the pivotal state (either Ohio or Florida). This is the assumption embodied in Equation (11).

The ranking model is the natural opposite of the independence model; as will become clear in the next section, the ranking model assumes that the evolution of uncertainty in a state regarding how it will vote on election day is perfectly rank correlated with the evolution of such uncertainty in every other state. This model makes intuitive sense if traders believe that voters haven't yet made up their minds, are responding to news, and are responding in a way that is rank preserving; for example, if state $i$ is more conservative than state $j$, then it will remain so regardless of what news is revealed.

We now propose a simple model that embeds both of the above models and allows the possibility that a candidate will win any pair of states to lie anywhere between being independent and perfectly correlated.

### 5.3. A Diffusion Model

Let $Z_i(t)$ be a latent "sentiment" variable for state $i$ at time $t$, where the election is to be held at time $\tau \geq t$. The latent variable is defined such that the candidate in question wins state $i$ if and only if the sentiment in state $i$ on election day is positive, that is, if $Z_i(\tau) > 0$.[5] To make this model tractable and allow estimation of its parameters, we assume that $Z_i(t)$ is a standard diffusion process

$$dZ_i(t) = dW_i(t) \tag{15}$$

for $i = 1, 2, \ldots, 51$ where $W_i(t)$ is a standard Wiener process. The assumption that the process is driftless can be seen as a consistency condition: it requires that the probability of a candidate winning any set of states on election day follows the law of iterated expectations. The assumption of a unit variance is simply a scaling property. The latent variable is itself unobservable, and because our candidate wins state $i$ if and only if $Z_i(\tau) > 0$, any process $Z_i'(t)$ with $\sigma_i \neq 1$ can be replaced by a scaled process $Z_i(t) = Z_i'(t)/\sigma_i$, which has the same level crossing properties at zero and does have a standard deviation of one. We further assume that

$$E[dW_i(t)dW_j(t)] = \rho_{ij}dt; \tag{16}$$

thus, state-specific shocks are correlated across states. Given these assumptions, it follows easily that

$$E[Z_i(\tau) \mid z_i(t)] = z_i(t), \tag{17}$$

$$\mathrm{Var}(Z_i(\tau) \mid z_i(t)) = \tau - t, \tag{18}$$

$$\mathrm{Cov}(Z_i(\tau), Z_j(\tau) \mid z_i(t), z_j(t)) = \rho_{ij}(\tau - t), \tag{19}$$

$$p_i(t) \equiv \Pr\{Z_i(\tau) > 0 \mid z_i(t)\} = \Phi\left(\frac{z_i(t)}{\sqrt{\tau - t}}\right), \tag{20}$$

[5] For concreteness, $Z_i(\tau)$ could be interpreted as the vote count differential on election day, but it need not have this interpretation.

**Figure 3 Florida and Ohio Securities**

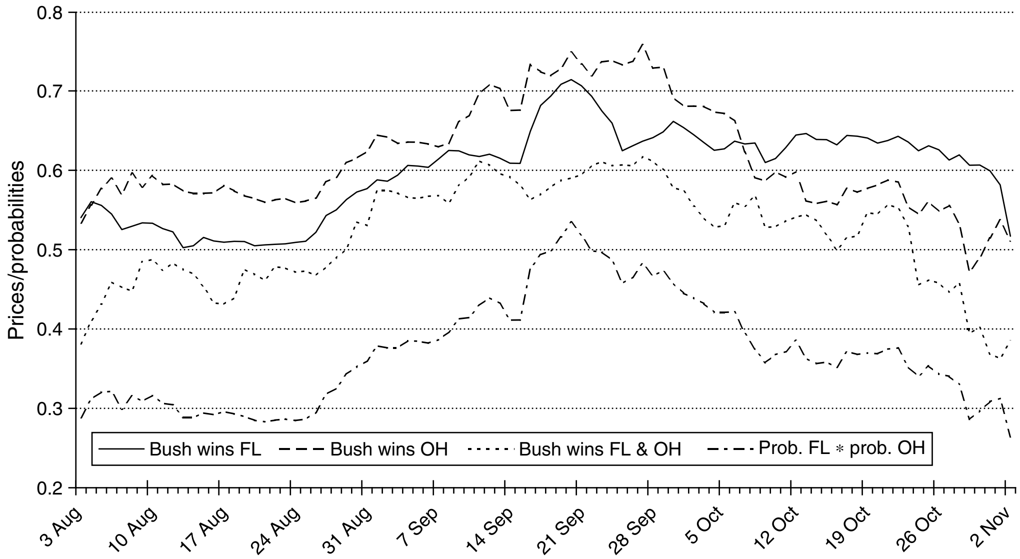


$$
\begin{aligned}
p_{ij}(t) &\equiv \Pr\{Z_i(\tau) > 0, Z_j(\tau) > 0 \mid z_i(t), z_j(t)\} \\
&= \Phi_2\left(\frac{z_i(t)}{\sqrt{\tau - t}}, \frac{z_j(t)}{\sqrt{\tau - t}}, \rho_{ij}\right) \\
&= \Phi_2(\Phi^{-1}(p_i(t)), \Phi^{-1}(p_j(t)), \rho_{ij}), \qquad (21)
\end{aligned}
$$

where $\Phi(\cdot)$ is the cumulative standard normal, $\Phi^{-1}(\cdot)$ is its inverse, $\Phi_2(\cdot, \cdot, \rho)$ is the bivariate cumulative standard normal with correlation $\rho$, and $z_i(t)$ is the sentiment realization in state $i$ at time $t$. Note that this diffusion model nests both the independence model (when $\rho = 0$ and hence $p_{ij}(t) = p_i(t)p_j(t)$ via Equation (21)) and the ranking model (when $\rho = 1$ and hence $p_{ij}(t) = \min(p_i(t), p_j(t))$ via Equation (21)).[6]

Intuitively, at each point in time, sentiment in each state is represented by a point on the real line that moves according to a driftless random walk, with steps drawn from the standard normal. The state will be won by the candidate in question if on the date of the election this point is positive; the state will be lost if the point is negative. The movement of each state can be taken to represent the reaction of that state to news that is being released as the election draws closer, and correlation between the movements of states represents the ways that these revelations affect voters (as reflected by traders) in correlated ways.

[6] The ranking model only requires that the sentiment variables stay ordered, not that they are perfectly correlated. However, because the sentiment variables are not observed but only latent in our model, perfect correlation can be assumed for the ranking model with no loss of generality.

## 6. Fitting the Diffusion Model

### 6.1. Direct Evidence of Correlation Between State Outcomes

As a first step, we now present direct evidence that some intermediate level of correlation between state outcomes was present in the Intrade data. To do this, we examine the prices for contracts that explicitly price the probability that Bush (or Kerry) will win some small set of states and compare it to the prices of those states in total. These contracts were present on Intrade for several regions (for example, Kerry to sweep the West Coast, Bush to sweep the South) as well as important "swing" states. The simplest example was the contract "Bush to win Florida and Ohio," which bundled together two large states that were widely regarded to be important to winning the election.

In Figure 3 we graph the daily prices of Bush to win: Florida, Ohio, Florida and Ohio, and the product of the first two probabilities. According to the independence model, Bush to win Florida and Ohio (the short-dashed curve) should be priced as the product of the prices for Bush to win each state individually (the bottom curve in the figure). According to the ranking model, the price for Bush to win Florida and Ohio should track the minimum of the two individual state prices (the minimum of the solid and long-dashed curves). Figure 3 shows clearly that the observed price for Bush to win both Florida and Ohio is strictly between the product of the two state prices and the minimum of the state prices, though it is closer to the minimum than it is to the product. This suggests that neither independence nor perfect

**Figure 4 Best Fit Rho for Florida and Ohio**

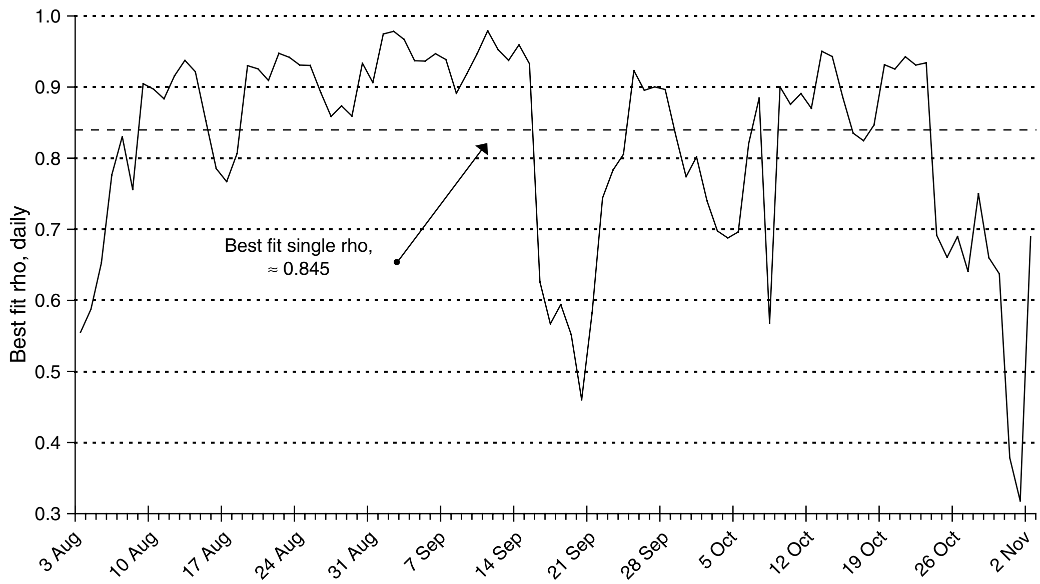


correlation adequately explains the pricing of the Florida and Ohio contract.

From Equation (21) we see that for every day $t$ and prices for "Bush to win Florida" and "Bush to win Ohio," there is a unique $\rho_{\mathrm{FL,OH}}(t)$ that will produce that day's observed price for "Bush to win Florida and Ohio."[7] In Figure 4 we plot the value of this $\rho_{\mathrm{FL,\,OH}}(t)$ over time. As expected, the resulting $\rho_{\mathrm{FL,\,OH}}(t)$s seem closer to one (the ranking model) than to zero (the independence model). The formal model developed in the previous section requires a constant correlation $\rho_{\mathrm{FL,\,OH}}$ over time, so we used least squares to find the correlation that best fits the whole time series. The resulting $\rho_{\mathrm{FL,\,OH}} \approx 0.85$, which confirms this intuition. In Figure 5 we plot the results of the diffusion model with a constant $\rho_{\mathrm{FL,\,OH}} = 0.85$ and compare these results to both the ranking and the independence model. Roughly speaking, the diffusion model matches the movements of the observed joint security quite well at the weekly level but seems to miss much of the variation at higher frequencies (as do both the ranking and independence models).

The diffusion model can also price securities that bundle election outcomes in more than two states via the multivariate extension of Equation (21). Consider a security for the event that a candidate wins in all states within some specified bundle $B$. The diffusion model would price such a security as

$$p_B(t) = \Pr\left\{\bigcap_{i\in B} Z_i(\tau) > 0 \,\middle|\, z_i(t)\ \forall\, i \in B\right\}, \tag{22}$$

where the probability is calculated from the multivariate normal distribution with means and covariances as specified in Equations (17)–(19). Pricing such bundles under the independence and ranking models is simple, for under independence we have

$$p_B(t) = \prod_{i\in B} p_i(t) \tag{23}$$

and the ranking model implies that

$$p_B(t) = \min_{i\in B} p_i(t). \tag{24}$$

Table 1 reports the results of using least squares to fit diffusion models to five different bundled securities that were traded over 100 days concluding with the 2004 election. Each diffusion was constrained to have the same (bundle-specific) pairwise sentiment correlation $\rho$ across all pairs of states, so only one free parameter was estimated per diffusion. We report the optimal correlation estimated for each bundle, along with the root mean squared error for the diffusion, independence, and ranking models; the root mean squared error can be interpreted as the magnitude in cents of the typical daily pricing errors made by the various models in replicating the observed prediction market prices.

Most obvious from Table 1 is the extent of sentiment dependence across states evidenced in these

[7] This would not be true, of course, were there any days with prices for "Bush to win Florida and Ohio" greater than the minimum of the two prices. There were no such days in the sample for Florida and Ohio. However, there were such days for other securities. For example, on September 20, 2004, a security for Kerry to win California, Oregon, and Washington traded at 68 cents and the security for Bush to win Oregon traded at 36 cents, implying a 64-cent valuation for Kerry. The chance that Kerry would win Oregon must be at least as large as the chance he would win Oregon and California and Washington, so these prices are inconsistent.

**Figure 5 Comparing Models on Florida and Ohio**

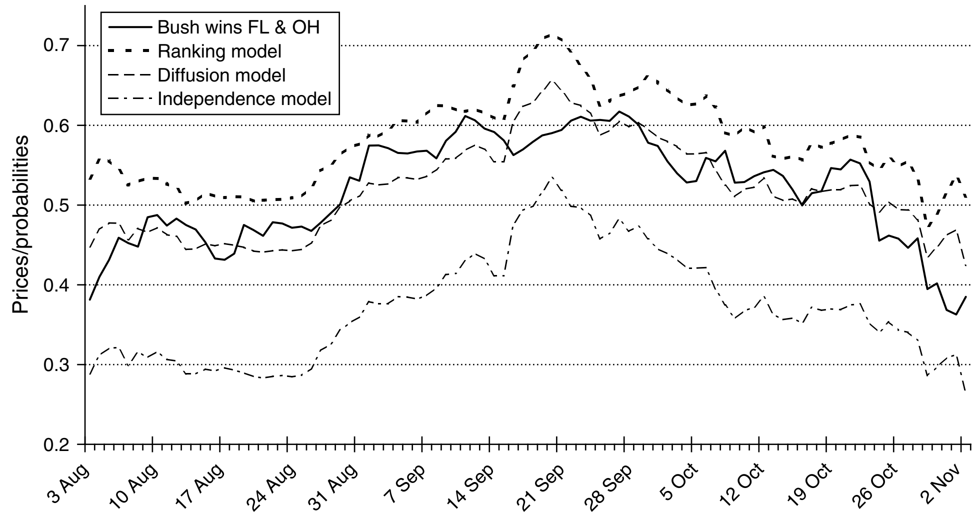


bundled securities via the estimated values of $\rho$. As was the case with Florida and Ohio, ignoring this dependence and assuming independence results in large pricing errors; indeed the largest root mean squared pricing error for the diffusion model (7.4 cents for Bush sweeps the South) is half the size of the smallest root mean squared pricing error for the independence model (14.7 cents for Bush sweeps the Southwest). The ranking model, which assumes $\rho = 1$ across all pairs of states, fairs better. For the security on Kerry to sweep the West Coast, the ranking and diffusion models are essentially the same (with the estimated sentiment correlation equal to 0.96), and both models report root mean square daily pricing errors of about 3.3 cents. For other bundles, however, the ranking model is less successful with pricing errors nearly twice those of the diffusion in some cases.

**Table 1 Modeling Bundled Security Prices**

| | | Root mean squared error | | |
|---|---|---|---|---|
| Bundle | Optimal $\rho$ | Diffusion | Independence | Ranking |
| Kerry sweeps West Coast[a] | 0.96 | 3.26 | 18.04 | 3.28 |
| Bush sweeps South[b] | 0.71 | 7.37 | 29.90 | 12.02 |
| Bush sweeps Great Plains[c] | 0.69 | 4.90 | 15.06 | 9.24 |
| Bush sweeps Southwest[d] | 0.65 | 3.56 | 14.72 | 6.84 |
| Kerry sweeps New England[e] | 0.63 | 4.54 | 15.70 | 7.87 |

[a]California, Oregon, Washington.
[b]South Carolina, Mississippi, Florida, Alabama, Georgia, Louisiana, Texas, Virginia, Arkansas, North Carolina, Tennessee.
[c]Iowa, Kansas, Minnesota, Nebraska, North Dakota, Oklahoma, South Dakota, Texas.
[d]Nevada, New Mexico, Utah, Colorado.
[e]Connecticut, Rhode Island, Maine, Vermont, Massachusetts, New Hampshire.

Focusing on the estimated values of $\rho$, it is interesting that all of these are relatively large, as was the sentiment correlation of 0.85 for Florida and Ohio. The bundled securities cover states that are geographically contiguous and traditionally viewed as voicing similar political concerns, so it is perhaps not surprising to discover that Intrade prices reflect highly correlated sentiment across these states. Florida and Ohio, though not contiguous, were key swing states where both the Republican and Democratic parties spent large sums of money campaigning during the election. It is perfectly reasonable to have a large correlation between these two states if they were swing states for similar reasons. When considering *all* states in the nation, however, should one still expect highly correlated sentiment across all states?

### 6.2. Fitting the Diffusion Model to Electoral Cutoff Prices

Turning to the national election, we consider "electoral cutoff" contracts of the form "Bush to win $\geq x$ electoral college votes." Letting $T$ denote the total number of electoral college votes garnered by Bush, the prices on such contracts provide point estimates of $\Pr\{T \geq x\}$. To extend the analysis of the previous section to these cutoff contracts, we fit our diffusion model to the observed cutoff security prices assuming that there is a single correlation $\rho$ responsible for the correlated sentiment shocks between any two states; that is, we assume that $\rho_{ij} = \rho$ in Equations (16) and (19). This permits us to seek the value of $\rho$ that leads to estimated security prices that best match the observed cutoff security prices for $x = 150, 200, 250, 269, 300, 350, 400$, and 450.

We use simulation to estimate the cutoff security prices that correspond to a given value of $\rho$. On any

**Figure 6 Rho Fitted to Daily Electoral-Vote Prices**

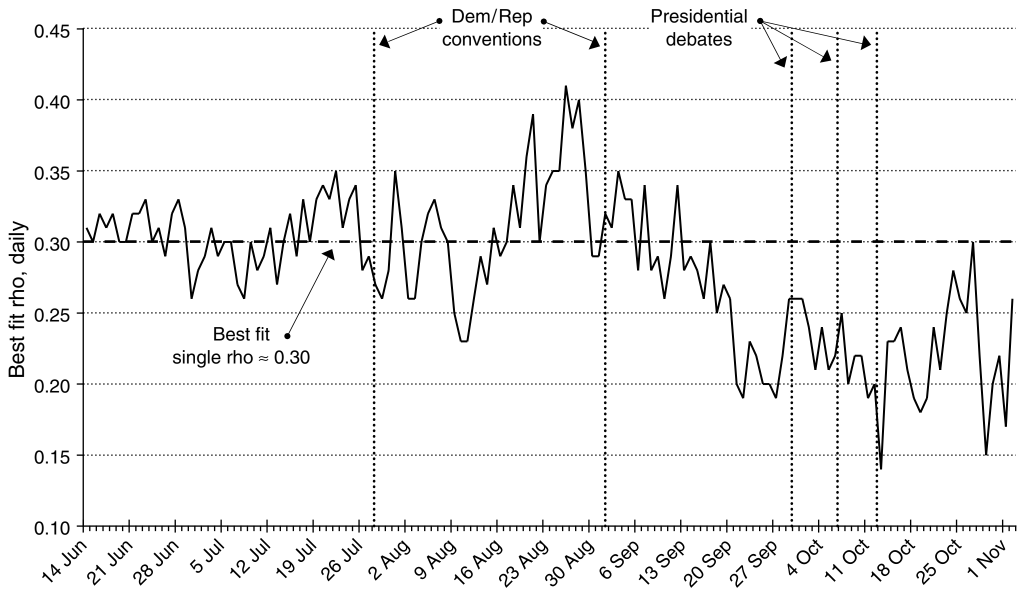


given day, we first obtain the sentiment value $z_i(t)$ for each state $i$ from the observed price $p_i(t)$ on the standard "Bush wins state $i$" contract by inverting Equation (20) to obtain

$$z_i(t) = \sqrt{\tau - t}\, \Phi^{-1}(p_i(t)) \quad i = 1, 2, \ldots, 51. \tag{25}$$

Following Equations (17), (18), and (19), the $r$th replication of the election day sentiment values $Z_1^r(\tau, \rho), Z_2^r(\tau, \rho), \ldots, Z_{51}^r(\tau, \rho)$ are simulated from a multivariate normal distribution with means $z_i(t)$, variances $\tau - t$, and pairwise correlation $\rho$ for *all* state pairs. The number of electoral college votes for Bush $T^r(\rho)$ corresponding to these sentiment values is then found by tallying the number of electoral votes over all states $i$ for which $Z_i^r(\tau, \rho) > 0$, that is,

$$T^r(\rho) = \sum_{Z_i^r(\tau, \rho) > 0} v_i, \quad r = 1, 2, \ldots, n, \tag{26}$$

where $n$ is the number of simulated replications. The estimated cutoff security price corresponding to Bush winning at least $x$ electoral college votes assuming a correlation of $\rho$, $\widehat{\Pr}\{T(\rho) \geq x\}$, is then given by

$$\widehat{\Pr}\{T(\rho) \geq x\} = \frac{1}{n} \sum_{r=1}^{n} 1_{\{T^r(\rho) \geq x\}}. \tag{27}$$

We evaluate the goodness of fit for a particular value of $\rho$ via the squared error criterion:

$$\sum_x (\Pr\{T \geq x\} - \widehat{\Pr}\{T(\rho) \geq x\})^2. \tag{28}$$

For each day we determine the value $\rho^*(t)$ that minimizes the sum of squares between the observed and modeled cutoff security prices using $n = 5{,}000$ simulated replications of election day sentiment for all values of $\rho(t)$ between 0 and 1 inclusive in increments of 0.01.

Figure 6 plots these values over time. Note that all values of $\rho^*(t)$ fall between 0.15 and 0.4, considerably lower than the estimates we obtained earlier for state-bundled securities. Though we are able to produce a different estimate $\rho^*(t)$ for each day $t$ of the election campaign, the formal theory underlying our diffusion model requires that $\rho$ remains constant over time (as do the independence and ranking models). We therefore searched for the best fit constant $\rho^{**}$ that minimizes the sum of squared differences between observed and estimated cutoff security prices over the entire time period available in the data. This produced a point estimate of $\rho^{**} \approx 0.30$. Recall that the diffusion model nests both the independence model ($\rho = 0$) and Fair's ranking model ($\rho = 1$). Thus, the distribution of total electoral votes (as approximated by the electoral cutoff security prices) is closer to the independence extreme than the ranking model.[8]

Why is $\rho^{**}$ so much smaller when estimated for the entire nation than for state-specific bundles? The answer, we believe, is that for many pairs of states,

[8] A reviewer noted correctly that some of the state-specific securities had low volume, resulting in wide bid-ask spreads, which could reduce the accuracy of midpoint pricing, impacting our estimation of $\rho$. However, as thin trading is a feature of quite certain (security prices near 0 or 1) rather than highly uncertain (security prices near 0.5) events, we conducted a robustness check by assigning all observed (midpoint) prices of at least 0.95 to unity, prices of at most 0.05 to 0, and repeated the analysis (thus removing any influence that thinly traded securities might have on the results). The best fitting correlation $\rho^{**}$ only changed from 0.30 to 0.33.

underlying sentiment correlation is close to zero if not exactly that. In some states, sentiment favoring Bush was perhaps nearly constant during the run up to the election (e.g., large and positive in "solid red" states and large and negative in "solid blue" states), implying nearly zero sentiment correlations between such states and *any* other state. Also, with $\binom{51}{2} = 1{,}275$ possible pairwise sentiment correlations, there are many more state pairings across rather than within contiguous regions, and one might expect sentiment movement to be stronger within than across regions. Alabama and Louisiana might react similarly to unfolding events prior to the election, but should one really expect similar association between Alabama and Massachusetts (or Oregon or Hawaii)?

As a check on the consistency of our national diffusion model, we reestimated $\rho$ based solely on the extreme contracts that Bush would win at least 350 electoral college votes or that Bush would win at least 200 electoral college votes, for as a reviewer suggested, one might expect that political shocks sufficient to give Bush at least 350 votes or fewer than 200 votes (which would price as 1 minus the price of Bush receiving at least 200 votes) could have stronger cross-state sentiment correlation. Of course, if the diffusion model is correct, we should obtain similar estimates of $\rho$ based only on these more extreme contracts to what we found earlier ($\rho^{**} \approx 0.30$), for swings in the prices of these extreme contracts would be consistent with the same diffusion process governing all contracts. Following the same procedures outlined above but basing goodness of fit solely on the squared prediction error resulting from estimating $\Pr\{T(\rho) \geq 350\}$ or $\Pr\{T(\rho) \geq 200\}$, respectively, we found that the optimal $\rho$ was 0.26 for the "Bush wins at least 350" contract and 0.30 for the cutoff at 200 electoral votes. Both of these are consistent with our earlier estimate of $\rho^{**} \approx 0.30$ based on the entire family of electoral college cutoff securities, providing further support for the diffusion model.

## 7. Model Comparisons

### 7.1. Variability of the Electoral College Distribution

Recall from Equation (4) the general formula for the variance of the number of electoral college votes. On any given day, we can estimate this directly from observed security prices under the assumptions of each of our three models: independence, ranking, and diffusion. The three models differ in their assessment of the joint probability $p_{ij}(t)$, which is the probability assessed on day $t$ that the candidate will carry both state $i$ and state $j$ on election day at time $\tau$. The independence model sets $p_{ij}(t) = p_i(t)p_j(t)$; the ranking model sets $p_{ij}(t) = \min(p_i(t), p_j(t))$, and the diffusion model sets $p_{ij}(t) = \Phi_2(\Phi^{-1}(p_i(t)), \Phi^{-1}(p_j(t)), \rho^{**})$.

To produce upper and lower bounds on the variance of the electoral college distribution implied by the Intrade data, we again use the following cutoff securities that traded during the election campaign: Bush to win at least 150, 200, 250, 269 (the election-winning security), 300, 350, 400, or 450 electoral college votes. Let $x_i$ denote these eight electoral college vote cutoffs for $i = 1, 2, \ldots, 8$, and define $x_0 = 0$ and $x_9 = 539$. These cutoffs define nine bins with occupancy probabilities given by $\omega_i$ where

$$\omega_i = \Pr\{T \geq x_{i-1}\} - \Pr\{T \geq x_i\} \quad \text{for } i = 1, 2, \ldots, 9. \tag{29}$$

Let $i_\mu$ denote the bin that contains the expected number of electoral college votes $\mu$ implied by the state-by-state "Bush wins" contracts. An upper bound for the variance of the electoral college distribution is then given by

$$\sigma^2_{\max} = \sum_{i=1}^{i_\mu - 1} (x_{i-1} - \mu)^2 \omega_i + \sum_{i=i_\mu+1}^{9} (x_i - \mu)^2 \omega_i + \max(x_{i_\mu} - \mu, x_{i_\mu - 1} - \mu)^2 \omega_{i_\mu}, \tag{30}$$

while a lower bound is given by

$$\sigma^2_{\min} = \sum_{i=1}^{i_\mu - 1} (x_i - \mu)^2 \omega_i + \sum_{i=i_\mu+1}^{9} (x_{i-1} - \mu)^2 \omega_i. \tag{31}$$

We employ Equations (30) and (31) to compute upper and lower bounds for the variance of the electoral college distribution for each day during the run up to the election.

In Figure 7, we plot the resulting standard deviations of the electoral college distribution for each model as assessed on each day in our sample along with standard deviations corresponding to the bounds presented above. Though the bounds are wide, they serve our purpose of judging the plausibility of the three models for the electoral college distribution under discussion. The standard deviations of the independence model are clearly too small, falling well below the lower bound derived from the cutoff security data for all days in our sample. Similarly, the standard deviations of the ranking model flirt with the upper bound derived from the cutoff security data over the entirety of our sample and often exceed this bound. Only the diffusion model sits comfortably between the data-derived upper and lower bounds.

Note from Figure 7 that the standard deviation of the number of electoral college votes declines as election day approaches, something we alluded to when discussing Figure 2. This feature is also

**Figure 7** **Standard Deviation of Electoral College Distribution**

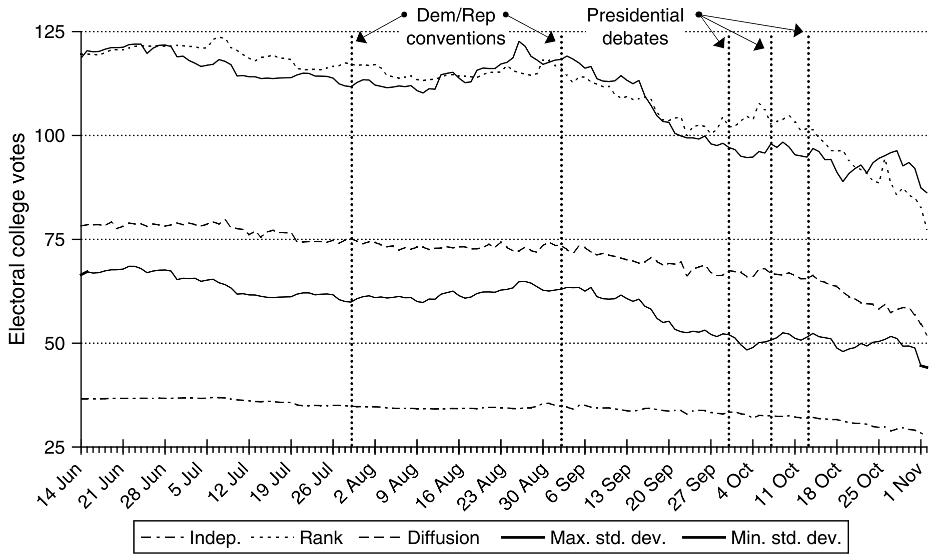


consistent with the diffusion model: Equation (18) shows that under the diffusion model, as election day approaches, the conditional variance of election-day sentiment dissipates in all states (as there is less time for new sentiment-shifting information to be revealed). This pushes $p_i(t)$, the probability of winning in state $i$ evaluated at time $t$, toward either zero or one (depending on whether the sentiment $z_i(t)$ is positive or negative), which in turn reduces the variability of the electoral college distribution.

### 7.2. Model Comparisons: Predicting the National Election

Now we compare the diffusion, independence, and ranking models on their ability to aggregate state-winning probabilities into the probability that Bush will win the presidential election. Recall that this probability can be determined for the independence model by calculating the exact probability distribution of $T$ from Equation (10); for the ranking model this probability is given by Equation (14). For the diffusion model, we refer to our simulations, discussed earlier, and estimate the probability that Bush will win the election using Equation (27); that is, we set

$$\Pr(\text{Bush Wins}) = \widehat{\Pr}\{T(\rho^{**}) \geq 269\} = \frac{1}{n}\sum_{r=1}^{n} 1_{\{T^r(\rho^{**})\geq 269\}}. \qquad (32)$$

In Figure 8, we plot the differences between the probabilities each model assigns to Bush winning the election, given each day's observed state prices and

**Figure 8** **Prediction Errors for the Probability Bush Wins**

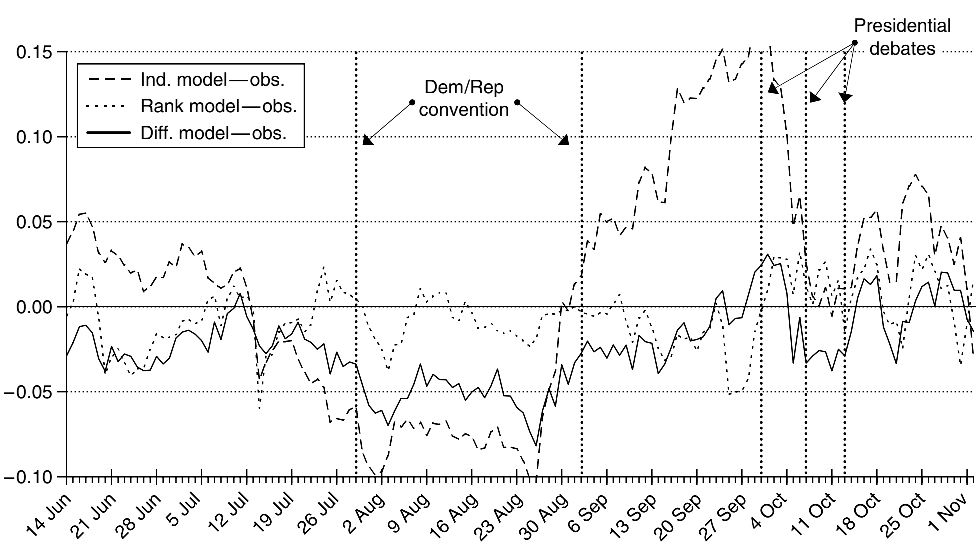

the observed price of the "Bush to win the U.S. presidency contract" traded on that day. Figure 8 shows clearly that both the ranking model and the diffusion model perform much better than the independence model at replicating the price of the "Bush win" contract over time. Note how the independence model's underestimation of the variability in the electoral college distribution leads it to overreact to upward and downward trends in the market. Because the probability of winning the presidency is a threshold (or "tail") probability, as the expected number of electoral college votes crosses 269, the concentrated probability mass near the mean for the independence model pushes the probability of receiving at least 269 votes much higher than what results from the relatively more diffuse distributions derived from the ranking and diffusion models (recall that all three models are constrained to produce the same expected number of electoral college votes as implied by the state-by-state "Bush win" contracts each day).

Focusing on the diffusion and ranking models, Figure 8 makes clear that although both models typically produce win probabilities within 5 cents of the trading price on the "Bush win" security, the ranking model provides a much better match to the data, especially during the period July 20 through mid-September. During this period, the diffusion model predicted that Bush would win the presidency with probability lower than 1/2, although both the ranking model and the observed price stay comfortably above 1/2. During this period, the expected number of electoral college votes for Bush, $E(T)$, consistently fell below 269, even though the price of the "Bush win" contract was trading above 50 cents, implying a win probability in excess of 50%. These two predictions can be consistent with each other, but only if the distribution of electoral votes for Bush has a thicker lower than upper tail (e.g., if the mean is less than the median). Our best fitting "constant rho" diffusion model with $\rho^{**} = 0.30$ does not produce such a distribution, so even though our diffusion model produces a better match to the overall electoral college distribution than the ranking model (which is also a diffusion model but with $\rho = 1$), perhaps surprisingly the ranking model does a better (in fact astonishingly good) job at matching the observed price on the presidential win security. Indeed, as Fair (2007) reports, the ranking model correctly identified all states that actually voted for or against President Bush while fingering Ohio as the pivotal state in the 2004 election (recall that Bush won the election with 286 electoral votes).

## 8. Conclusion

With prediction markets gaining popularity, it is important to understand their behavior. The presidential election provides a useful test case, for to be useful, the pricing of the various securities traded must be internally consistent. We began by establishing simple probabilistic relationships among the key securities traded in this market, namely, the prices on the state-by-state "Bush wins" contracts and the prices on electoral cutoff securities of the form "Bush wins at least $x$ electoral college votes," which includes the election-winning security of garnering at least 269 electoral college votes. We then showed that the observed prices for the state-by-state contracts are reasonably consistent with the prices on the electoral-cutoff securities by interpreting the latter as points along the survivor distribution of the total number of electoral college votes.

After reporting two benchmark models already in the literature—Kaplan and Barnett's independence model and Fair's ranking model—we developed a new diffusion model that nests both of these. The diffusion model suggests in a simple way how the prediction market incorporates new information over time to revise the estimated probabilities of Bush winning each state. We fit this model to the electoral college distribution implied by the electoral-cutoff security prices and discovered that the underlying "sentiments" implied by the market data were correlated across different states.

Our estimate of $\rho = 0.30$ stands in contrast to the extreme correlations of zero implied by the independence model and one implied by the ranking model. Indeed, we showed that the variance of the electoral college distribution resulting from the independence model always fell below a lower bound based solely on the market data during the run up to the election, and the ranking model produced a variance that flirted with the upper bound over this same time period. Our best fitting diffusion model produced a variance that fell comfortably within the upper and lower bounds throughout the entire sample period. With regard to pricing the presidential win security, we showed that as a consequence of ignoring the correlation in sentiment across states, the independence model overreacted to upward and downward swings in the market-implied expected number of electoral college votes for Bush, which in turn led to over- and underestimates of the probability of Bush winning the election over time. The ranking and best fitting diffusion models provided much better matches to the observed price on the presidential win security, though perhaps surprisingly the ranking model provided a much closer match to the observed price in spite of the fact that the diffusion model provided a better fit to the overall electoral college distribution.

One could certainly provide improvements to the models attempted here—for example, one could return to the diffusion model but allow a more general correlation structure to capture differentially related changes in sentiment across different pairs of states, perhaps with the aid of covariates. Nonetheless, although it is one thing to ask whether the price of an individual security reflects the underlying probability of the underlying contracted event (Manski 2006, Wolfers and Zitzewitz 2007), it is more challenging for the prices of each of a family of individually traded securities to jointly obey a common probability model consistent with rules such as those imposed by the electoral college. Our analysis thus leaves us with the overall impression that Intrade's 2004 presidential prediction market was remarkably consistent.

The diffusion model developed here can also be used to price more complex contracts that could be traded in the future. For example, we could imagine an option-like contract that paid not a fixed amount if one candidate won more than 280 electoral votes, but rather an amount proportional to the difference between the actual electoral votes received and 280; e.g., $\$10 \times \max(\text{Electoral votes} - 280, 0)$. To model the value of such a call option requires knowledge of the entire probability distribution of the electoral college vote. The methods developed in this paper allow us to do precisely that.

## Acknowledgments

The authors acknowledge Dominic Soon for invaluable research assistance, Ray Fair for his helpful comments on earlier versions of this work, and the associate editor and two anonymous reviewers for suggestions that greatly improved this paper. Comments are welcome at keith.chen@yale.edu or at 135 Prospect St., Box 208200, New Haven, CT 06520.

## References

Banz, R. W., M. H. Miller. 1978. Prices for state contingent claims: Some estimates and applications. *J. Bus.* **51**(4) 653–672.

Bodurtha, J. N., Q. Shen. 1999. Historical and implied measures of "value at risk": The DM and yen case. Working paper, Georgetown University, Washington, D.C.

Breeden, D., R. Litzenberger. 1978. Prices of state contingent claims implicit in option prices. *J. Bus.* **51**(4) 621–651.

Campa, J., K. Chang. 1998. The forecasting ability of correlations implied in foreign exchange options. *J. Internat. Money Finance* **17**(6) 855–880.

Erikson, R. S., K. Sigman. 2000. A simple stochastic model for close U.S. presidential elections. Working paper, Department of Political Science, Columbia University, New York.

Fair, R. C. 2002. *Predicting Presidential Elections and Other Things.* Stanford University Press, Stanford, CA.

Fair, R. C. 2007. Interpreting the predictive uncertainty of elections. Working paper, Cowles Foundation and Department of Economics, Yale University, New Haven, CT.

Jackman, S., D. Rivers. 2001. State-level election forecasting during election 2000 via dynamic Bayesian hierarchical modeling. Working paper, Department of Political Science, Stanford University, Stanford, CA.

Kaplan, E. H., A. Barnett. 2003. A new approach to estimating the probability of winning the presidency. *Oper. Res.* **51**(1) 32–40.

Latane, H., R. Rendleman. 1976. Standard deviations of stock price ratios implied in option prices. *J. Finance* **31**(2) 369–381.

Manski, C. F. 2006. Interpreting the predictions of prediction markets. *Econom. Lett.* **91** 425–429.

Skintzi, V., A. Refenes. 2005. Implied correlation index: A new measure of diversification. *J. Futures Markets* **25**(2) 171–197.

Wlezien, C., R. S. Erikson. 2007. The horse race: What polls reveal as the election campaign unfolds. *Internat. J. Public Opinion Res.* **19**(1) 74–88.

Wolfers, J., E. Zitzewitz. 2004. Prediction markets. *J. Econom. Perspectives* **18**(2) 107–126.

Wolfers, J., E. Zitzewitz. 2007. Interpreting prediction market prices as probabilities. Working paper, University of Pennsylvania, Philadelphia. http://bpp.wharton.upenn.edu/jwolfers/Papers/InterpretingPredictionMarketPrices.pdf.