Cowles Foundation
for Research in Economics at Yale University

Cowles Foundation Discussion Paper No. 1579

September 2006
Revised May 2007

INTERPRETING THE PREDICTIVE UNCERTAINTY OF ELECTIONS

Ray C. Fair

Interpreting the Predictive Uncertainty of Elections

Ray C. Fair*

Revised May 2007

Abstract

This paper provides an interpretation of the uncertainty that exists at the beginning of the day of an election as to who will win. It is based on the theory that there are a number of possible conditions of nature that can exist on election day, of which one is drawn. Political betting markets like Intrade provide a way of trying to estimate this uncertainty. It is argued that polling standard errors do not provide estimates of this type of uncertainty. They instead estimate sample-size uncertainty, which can be driven close to zero with a large enough sample.

This paper also introduces a ranking assumption concerning dependencies across U.S. states, which puts restrictions on the possible conditions of nature than can exist on election day. The joint hypothesis that the last-day Intrade ranking is correct and the ranking assumption is correct predicts the exact outcomes of the 2004 presidential election and the 2006 Senate election. Although not a test of the ranking assumption, there is evidence that the Intrade traders used the ranking assumption to price contracts in the 2004 presidential election. This was not the case, however, in the 2006 Senate election. Finally, it is shown if the ranking assumption is correct, the two political parties should spend all their money on a few states, which seems consistent with their actual behavior in 2004.

*Cowles Foundation and International Center for Finance, Yale University, New Haven, CT 06520-8281. Voice: 203-432-3715; Fax: 203-432-6167; email: ray.fair@yale.edu; website: http://fairmodel.econ.yale.edu.

I am indebted to Keith Chen, Don Green, Ed Kaplan, Steve Morris, Sharon Oster, Justin Wolfers, and Eric Zitzewitz for helpful discussions and comments and to Mike Knesevitch at Intrade for supplying me with some of the data. This paper is a substantially revised version of an earlier paper entitled "Predicting Electoral College Victory Probabilities from State Probability Data," November 2004.

1 Introduction

At the beginning of the day of an election, after all the campaigning has been completed, there is still uncertainty as to who will win. This paper provides an interpretation of what this uncertainty is. The theory is that there are a number of possible "conditions" of nature that can exist on election day, of which one is drawn on election day. The uncertainty is which condition will be drawn. Section 2 presents this theory. Section 3 then discusses a way in which this uncertainty can be estimated using political betting markets. Section 3 also discusses the use of polling standard errors to estimate uncertainty. It is argued that polling standard errors do not provide estimates of the type of uncertainty considered in this paper. The rest of the paper is concerned with an assumption about dependencies across U.S. states, called the "ranking" assumption, that puts restrictions on the possible conditions of nature than can exist on election day. This assumption is discussed and examined empirically in Section 4. Section 5 is concerned with the question of how the two political parties should behave regarding campaign spending across states in U.S. presidential elections if the ranking assumption is correct.

2 The Theory

It is assumed that on election day there are n possible conditions of nature regarding the events that will take place during the day, each with probability $1/n$ of occurring. If in p percent of the n conditions candidate A wins, then p is the probability that A wins. The theory is that there are many possible conditions left at the end of a campaign and that the uncertainty is which condition will be drawn. Even if one knew the n possible conditions of nature, the best that one could say at the beginning of election day is that A would win with probability p.

An alternative way of thinking about the beginning of election day is that there is only one condition of nature left. In this case if one knew the condition, the outcome would be known with certainty. There are many reasons, however, to think that there is more than one possible condition left. The weather is not known with certainty, and weather may affect turnout, which may benefit one candidate relative to the other. A voter's decision may be affected by events that happens to him or her a few hours before he or she enters the voting booth, and there may be more than one set of possible events. For example, in condition of nature 1 a voter driving to vote may glance at a sign that affects his vote, whereas in condition of nature 2 the voter does not see the sign because someone is honking behind him. Perhaps in one condition of nature a voter runs into a friend on the way to vote and the friend convinces her to vote for candidate A, whereas in another condition of nature she does not run into the friend and does not vote for A. Implicit in this theory is the view that people's voting decisions are affected by what happens to them during the day and that there are a number of possible things that can happen to them. Also, people's feelings and moods may vary from day to day, and so there may be a number of possible ways they can feel on election day. In short, the die has not been cast at the time the election begins.

Let $\mathcal{S}$ denote the set of possible conditions of nature on election day. During the course of a campaign, many things happen that can affect $\mathcal{S}$. If, for example, candidate A does poorly in a debate, this may eliminate a number of possible conditions of nature that otherwise would have existed on election day in which A wins. $\mathcal{S}$ is thus different than it otherwise would have been had A not done poorly in the debate. One can think of a campaign as trying to eliminate as many possible conditions of nature as possible in which its candidate loses.

It may be the case that $\mathcal{S}$ contains almost no possible conditions of nature in which one candidate loses a particular state. For example, in the 2004 U.S. presidential election, there were probably very few possible conditions of nature in which President Bush lost Texas. $\mathcal{S}$ could still be a large set, but a set in which almost all conditions contain a Bush win in Texas.

3 Estimating Uncertainty

Political Betting Markets

Political betting markets provide one way of trying to estimate the uncertainty just discussed. The market considered in this paper is Intrade.[1] Prior to the 2004 U.S. presidential election the website *www.intrade.com* allowed one to buy and sell contracts for each state and the District of Columbia. The contract for Iowa, for example, stated "G W Bush to win the electoral votes of Iowa." The contracts were in units of ten dollars, and a price of 55.0 meant that you could buy one contract for $5.50. If Bush won Iowa, you would get back $10.00. Otherwise, you would get back nothing. You could also sell the contract, winning $5.50 if Bush lost and losing $4.50 if Bush won. There was also a national contract that stated "George

[1] The Intrade data are sometimes referred to as Tradesports data. Intrade is a subdivision of Tradesports, and the data are the same.

W Bush is re-elected as United States President." There were also contracts for various combinations of state victories. For example, there was a Bush Greatplains contract that stated "Pres George W Bush to win IA, KS, MN, NE, ND, OK, SD, & TX." The national contract was by far the most traded contract on Intrade. The markets for many of the state contracts were fairly thin. An interesting discussion of this market and others like it is in Wolfers and Zitzewitz (2004a).

Let p_i denote the percent of the possible conditions of nature in which Bush wins state i. One possible estimate of p_i is the Intrade price that existed at the beginning of election day.[2] For example, a price of, say, 53.9 for Florida is interpreted as saying that the market expects that in 53.9 percent of the possible conditions of nature on election day Bush wins Florida.

The fifth column in Table 1 presents the prices of the state contracts that existed at 6:00 am Eastern time on the day of the election, November 2, 2004. This is the time at which the first poll opened. (Ignore for now columns one through four in Table 1.) The states are ranked in Table 1 by the prices in the fifth column, i.e., the prices on the last day. Many of the states have prices close to 100.0, and many have prices close to 0.0. A price close to 100.0 in the present context means that the market expects that there are very few possible conditions of nature in which Bush loses the state. The opposite is the case for prices close to 0.0.

Regarding the prices in Table 1, if one excludes the top 25 states through Missouri, which has a price of 87.1, and the bottom 15 states beginning with

[2] Manski (2004) has shown that under certain assumptions about the beliefs of traders the market price of a contract is not necessarily the mean belief of the traders. However, under what appear to be plausible assumptions, this bias is either zero or small—see Wolfers and Zitzewitz (2004b). This paper is based on the assumption that the bias is zero.

Table 1
Intrade Data

State	9/7	9/21	10/5	10/19	11/2	# Votes	$\sum$ Votes
	Intrade Price						
Montana	95.0	94.0	95.0	96.3	99.0	3	3
Oklahoma	97.0	97.0	97.0	97.0	98.4	7	10
Utah	96.0	97.0	97.0	97.5	98.0	5	15
Idaho	95.5	96.0	95.0	95.5	98.0	4	19
Texas	98.0	98.0	98.0	98.0	97.9	34	53
Wyoming	97.0	97.0	97.0	97.5	97.6	3	56
Indiana	96.0	96.0	91.2	94.4	97.4	11	67
Alaska	96.0	96.0	98.0	95.5	97.4	3	70
Louisiana	92.5	91.9	92.0	92.6	97.0	9	79
Tennessee	78.7	85.0	89.0	92.0	96.5	11	90
Kentucky	92.5	92.0	92.0	93.1	95.8	8	98
Kansas	96.0	96.0	93.5	94.1	95.8	6	104
Mississippi	96.0	96.0	94.0	94.5	95.6	6	110
Georgia	96.5	97.0	92.2	95.7	95.2	15	125
Alabama	98.0	96.0	94.0	96.5	95.2	9	134
Nebraska	96.0	97.5	94.0	95.7	95.2	5	139
South Carolina	95.0	97.0	91.0	93.7	95.1	8	147
North Dakota	96.0	96.0	92.5	95.5	95.1	3	150
South Dakota	96.0	96.0	92.0	95.7	95.1	3	153
North Carolina	81.0	93.0	87.5	89.0	94.7	15	168
Arizona	78.0	83.0	83.0	90.0	94.0	10	178
Virginia	86.0	91.0	87.5	87.8	93.2	13	191
West Virginia	67.7	77.0	77.0	79.9	92.0	5	196
Arkansas	73.0	78.0	84.0	82.0	90.0	6	202
Missouri	67.0	85.0	84.0	81.0	87.1	11	213
Colorado	75.5	76.0	75.0	79.4	77.0	9	222
Nevada	60.0	69.9	74.5	67.5	76.8	5	227
New Mexico	43.0	40.0	37.7	37.2	56.5	5	232
Florida	60.5	70.0	63.5	66.0	53.9	27	259
Ohio	63.0	72.0	67.5	57.8	51.1	20	279
Iowa	43.0	55.0	57.0	55.2	51.0	7	
Wisconsin	57.0	62.0	64.0	54.5	41.0	10	
New Hampshire	42.0	55.0	51.0	43.0	31.0	4	
Pennsylvania	43.4	43.0	35.0	38.0	28.9	21	
Hawaii	10.0	10.0	8.0	5.5	26.1	4	
Minnesota	40.0	40.5	35.5	38.5	24.0	10	
Michigan	33.0	29.9	23.0	19.9	11.1	17	
New Jersey	15.9	24.0	18.0	16.5	10.0	15	
Oregon	36.3	35.0	26.9	21.9	10.0	7	
Maine	27.4	26.2	26.5	24.0	9.2	4	
Delaware	16.0	18.0	13.0	9.6	5.1	3	
California	9.6	11.4	8.0	6.0	3.3	55	
Connecticut	8.0	7.0	7.0	5.7	3.3	7	
Washington	28.0	25.0	19.0	8.0	3.0	11	
Vermont	7.0	8.0	8.0	3.3	2.5	3	
Illinois	8.8	12.0	8.8	6.8	2.0	21	
Maryland	14.0	16.0	17.9	9.0	2.0	10	
New York	7.0	9.9	8.4	4.9	1.7	31	
Massachusetts	3.0	4.0	2.0	2.8	1.7	12	
Rhode Island	4.0	4.0	4.0	3.5	1.7	4	
DC	1.0	1.0	2.0	1.5	0.8	3	

- Votes are electoral votes. 269 votes are needed to win for President Bush.
- President Bush won Iowa, all the states above it, and none below it.

Michigan, which has a price of 11.1, there are 11 states left, ranging from Minnesota with a price of 24.0 to Colorado with a price of 77.0. The 3 closest states are Florida (53.9), Ohio (51.1), and Iowa (51.0). The prices from these 11 states are analyzed below.

Polling Standard Errors

Although polling standard errors are commonly used to estimate uncertainty, this type of uncertainty is not the same as the type that is of concern in this paper. In other words, the type of uncertainty discussed in Section 2 is not the type estimated by polling standard errors. Almost all polling organizations release both a mean prediction and a standard error of the mean prediction, and these standard errors estimate sample-size uncertainty. The larger the sample, the smaller the standard error. To see why this uncertainty is different from that discussed in Section 2, consider the extreme case in which every eligible voter were asked the day before the election whether he or she was planning to vote and for whom. This would yield a mean vote share with a standard error of zero.[3] On this score, there would be no uncertainty left, whereas the uncertainty discussed in Section 2 would still exist.

To examine uncertainty estimates from polling standard errors versus those from political betting markets, one can compare the probability of winning a state that is backed out from state polling data with the probability as estimated by the

[3]Even if the sample size were, say, only 100,000 eligible voters rather than all eligible voters, the standard error would be close to zero. For a binomial distribution with p equal to .5 and N equal to 100,000, the standard error of the mean is .0016.

Intrade price. Table 2 makes this comparison for the 11 states mentioned above except Hawaii. The polling data are from the Real Clear Politics (RCP) website.[4] Results for the Zogby poll were used along with the RCP average of a number of polls. The last date of the polls was November 1. (Hawaii was not used because the last date of a poll for it was October 20.) The sample size for each Zogby state poll was 601 likely voters. Zogby reported its standard error as 2.05 percent for each state, which is consistent with the sample size of 601 for a binomial distribution. The state standard errors on the RCP website varied for the different polls from about 1.5 to 2.5 percent, with the sample sizes varying from about 500 to 1,500. The backed out probabilities for Zogby in Table 2 are based the assumption of a normal distribution and a standard error of 2.05 percent. No standard errors were reported for the RCP average, and two choices are used in Table 2, 2.05 percent and 1.0 percent. The 1.0 standard error is consistent with a sample size of about 3,000. If the RCP average is an average of five polls, each with a sample size of 600, the total sample size is 3,000. Table 2 also lists for each state for Zogby and for the RCP average the estimated two-party vote share for Bush.

If the Intrade prices are picking up uncertainty not accounted for in the polling standard errors, i.e., the type of uncertainty discussed in Section 2, then one should expect for large sample sizes that the probabilities backed out of the polling data to be closer to either 0 or 100 than are the Intrade prices. Large sample sizes imply small polling standard errors and thus backed-out probabilities that are likely to be close to 0 or 100. In the present case it is unclear whether a standard error of 2.05 percent or even 1.0 percent is small enough for this property to hold, but it turns

[4] *http://www.realclearpolitics.com/bush_vs_kerry_sbys.html.*

Table 2
Intrade Prices versus Polling Data

State	Intrade Price	Backed out Probability Zogby[a]	Backed out Probability RCP[a]	Backed out Probability RCP[b]	Bush Share Zogby	Bush Share RCP
Colorado	77.0	68.7	90.6	99.6	51.0	52.7
Nevada	76.8	89.8	94.6	99.9	52.6	53.3
New Mexico	56.5	23.2	63.4	75.8	48.5	50.7
Florida	53.9	50.0	55.8	61.8	50.0	50.3
Ohio	51.1	94.6	70.4	86.4	53.3	51.1
Iowa	51.0	10.2	53.9	57.9	47.4	50.2
Wisconsin	41.0	6.5	59.6	69.1	46.9	50.5
New Hampshire	31.0	NA	40.4	30.8	NA	49.5
Pennsylvania	28.9	15.3	40.4	30.8	47.9	49.5
Minnesota	24.0	6.5	20.4	4.5	46.9	48.3

[a] Backed out probability based on a standard error of 2.05.
[a] Backed out probability based on a standard error of 1.00.
• RCP is the Real Clear Politics average of a number of polls.

out that it does hold for all but 4 of the 19 cases in Table 2 that use a standard error of 2.05 percent and for all but one of the 10 cases that use a standard error of 1.0 percent. For example, the Ohio Intrade price is 51.1, while the polling probabilities are 94.6, 70.4, and 86.4. For Wisconsin the Intrade price is 41.0 and the probabilities are 6.5, 59.6, and 69.1. In this case Zogby and RCP disagreed as to who would win, but both were more confident than Intrade. The 4 exceptions that use a standard error of 2.05 are Zogby Colorado (77.0 versus 68.7), Zogby Florida (53.9 versus 50.0), RCP New Hampshire (31.0 versus 40.4), and RCP Pennsylvania (28.9 versus 40.4). The one exception that uses a standard error of 1.0 is RCP Pennsylvania (28.9 versus 30.8).

An interesting example using backed-out probabilities from polling data is in

Leigh and Wolfers (2006). They find using data from the 2004 Australian election that they need to increase the polling standard errors to about 10 percent to get backed-out probabilities that are close to those from political betting markets. (A standard error of 10 percent is equivalent to a poll of only 25 voters.) They suggest (p. 334) from these results that pollsters' published margins of error should at least be doubled. The interpretation in the present paper, however, is simply that pollsters are estimating a different type of uncertainty. They are estimating sample-size uncertainty, whereas political betting markets are estimating the type of uncertainty discussed in Section 2. Even there is no sample-size uncertainty, there is still uncertainty.

If there were no uncertainty at the beginning of election day except sample-size uncertainty, then the Intrade prices would just be picking up the uncertainty reflected in the polling standard errors, i.e., in the polling sample sizes. In this case the Intrade prices would also approach 0 or 100 as the sample sizes increase. The results in Table 2, however, do not support this hypothesis. As noted above, most of the backed out probabilities from the polls are closer to 0 or 100 than are the Intrade prices, and generally they are quite different from the Intrade prices. This suggests that the backed-out probabilities and the Intrade prices are estimating different things, which is the argument of this paper. The results for the 2004 Australian election in Leigh and Wolfers (2006) also support this view.

Regarding the use of Intrade prices to estimate the type of uncertainty discussed in Section 2, one cannot rule out the possibility that these prices are in part affected by polling standard errors. So part of the uncertainty reflected in the Intrade prices might be sample-size uncertainty. If this is true, this bias will fall as the sample sizes increase, but if there is uncertainty of the type discussed in Section 2, the Intrade prices will not approach 0 or 100 as the sample sizes increase.

4 The Ranking Assumption: A Restriction on the Possible Conditions of Nature

The Ranking Assumption

The rest of this paper is concerned with an assumption about dependencies across U.S. states. It uses the "conditions of nature" framework in Section 2. The assumption, called the "ranking" assumption, puts restrictions on the possible conditions of nature than can exist on election day.

The assumption is easy to describe. Rank the states by p_i, as is done in Table 1 using the Intrade data. The assumption is then that there is no condition of nature in which Bush wins state i and loses a state ranked higher than i. If, for example, Texas is ranked higher than Massachusetts, then in none of the n conditions of nature does Bush win Massachusetts and lose Texas. There may be conditions in which Bush wins Massachusetts (Kerry makes some serious error), but in these conditions Bush also wins Texas.[5]

It is common in previous work to assume some form of independence. Kaplan and Barnett (2003) assume that the state outcomes are independent, that "the events that the candidate is leading in various states are mutually independent" (p. 33). Snyder (1989) analyzes districts and assumes that the elections in the districts are

[5] Ed Kaplan has pointed out to me that given a ranking like in Table 1, under the ranking assumption there are only 52 possible outcomes: Bush takes all 51, Bush takes all but the last one, Bush takes all but the last two, etc. This compares to 2^{51} possible outcomes, about 2.25 million billion. A remarkable economy of outcomes has been achieved by the ranking assumption!

all statistically independent. He points out that this rules out "uncertainty about national variables that may affect the electoral outcomes in all districts simultaneously, such as changes in aggregate output or foreign policy crises" (p. 646). Brams and Davis (1974) assume that "the voting of uncommitted voters within each state is statistically independent" (p. 120). Strömberg (2002) assumes that the state level popularity parameters of a candidate are independent, although he also has a national popularity parameter. Soumbatiants, Chappell, and Johnson (2006) have both national and state-specific shocks.

What would it mean in the present context for the state probabilities to be independent? On election day the probability of Bush winning state i is simply the percent of his state i wins in the n possible conditions of nature. The probabilities will, of course, change if the n possible conditions of nature change. Consider as a thought experiment different *sets* of n possible conditions of nature on election day. Say that Bush has done poorly in the debates in set 1 and well in set 2. One would expect all the state probabilities to be higher for Bush in set 2. In set 2 there would fewer conditions of nature in which Bush loses any given state. The state probabilities in this case would be positively correlated. In order for the probabilities to be uncorrelated, the sets must differ in state-specific ways. For example, the Republican party might be better organized in California in set 1 than in set 2, but everything else the same. The two sets would then differ only regarding the probability for California. These state-specific differences across different sets of the n possible conditions of nature seem less likely to occur than differences that affect all the state probabilities.

The ranking assumption does not, of course, directly concern different sets of the n possible conditions of nature. It simply puts restrictions on the n possible conditions of nature that exist on election day. If state i is ranked ahead of state j, then in no condition of nature does Bush win j and lose i. The concept of different *sets* of the n possible conditions of nature is not needed.

Tests of the Ranking Assumption using Intrade Data

After the outcome of the 2004 election the joint hypothesis that 1) the Intrade price ranking on the last day is correct and 2) the ranking assumption is correct can be tested. Under this joint hypothesis President Bush should not have won any state ranked below a state that he lost. Table 1 shows that he did not win any such state. Bush won Iowa, all the states above Iowa, and none below Iowa. The actual results are exactly as the joint hypothesis predicted.

Note from Table 1 that Bush won all the states with a price above 50 on the last day and lost all the states with a price below 50. Although this is obviously a plus for Intrade, it is not necessary for the joint hypothesis to be true. If, say, all the prices on the last day were 10 percent lower, so that the price of Iowa were 45.9 rather than 51.0, the results would still have been exactly as the joint hypothesis predicted even though Bush would have won Iowa with a price below 50.

Another test of the joint hypothesis can be made using Intrade data for the 2006 U.S. Senate election. Table 3 presents the last transaction price and the average of the bid and ask prices as of 6:00 AM on election day, November 7, 2006, for the seven states that were at all in play. Even for these seven states trading was thin, which is the reason for presenting both the last transaction price and the average

Table 3
Intrade Prices at 6:00 AM on November 7, 2006
The Prices are for a Democratic Win for the Senate

	State	Last	Average of Bid & Ask	Bid	Ask
1	New Jersey	87.8	92.3	90.5	94.0
2	Montana	79.9	78.8	75.0	82.6
3	Rhode Island	76.0	65.8	63.5	68.0
4	Maryland	70.2	72.5	71.0	74.0
5	Missouri	56.9	59.4	57.2	61.5
6	Virginia	53.5	55.0	54.0	56.0
7	Tennessee	12.0	15.0	12.0	18.0

• The Democrats won all but Tennessee.

of the bid and ask prices. The separate bid and ask prices are also presented to get a sense of the market. The states are ranked in the table by the last transaction price.

The ranking assumption says that the Democrats should not lose any state ranked above a state they won, and this was the case for the 2006 Senate races. The lowest ranked state that they won was Virginia (Virginia ranks lowest using either measure), and they won everything above Virginia. So as was the case for the 2004 presidential election, the 2006 Senate results are exactly as the joint hypothesis that the last-day Intrade ranking is correct and the ranking assumption is correct predicted.

It is also the case that Intrade is perfect for the Senate races in that the Democrats won every state with a price above 50 and no state with a price below 50. As noted above, Intrade does not have to be perfect in order for the ranking assumption to be perfect. For example, if the Democrats had lost Virginia, contrary to Intrade's prediction, the results would still have been exactly as the joint hypothesis predicted.

Do Intrade Traders Use the Ranking Assumption?

Unlike the above comparisons to the actual outcomes, the following is not a test of the ranking assumption. It is instead an examination of whether the Intrade traders are using the ranking assumption to price various contracts.

For this examination four other days of Intrade prices were sampled. These are presented in the first four columns in Table 1. The first date is September 7, 2004, the day after Labor Day. The other three are two weeks apart. The time of day is 10:00 am Eastern for the first, third, and fourth and 11:00 am Eastern for the second.

First, note that under the ranking assumption it is trivial to compute, given the individual state prices for any particular day, the probability that Bush wins in the Electoral College. Rank the states as is done in Table 1 for the last day and then go down the ranking, adding electoral votes, until 269 is reached. If this is state j, then state j is "pivotal," and the probability that Bush wins the election is simply the probability that he wins state j.

Now, given the individual state prices in Table 1, it turns out that the Intrade prices of various combination contracts are quite close to what one would expect if traders were using the ranking assumption. This can be seen in Table 4, which presents prices for various combination contracts along with what the ranking assumption would predict the prices should be and what the independence assumption would predict. For the Bush Greatplains contract, for example, the price predicted by the ranking assumption is the price of the lowest ranked state in the contract, which for September 7 is Minnesota with a price of 40.0. The price predicted by

Table 4
Intrade Prices for Various Contracts

	September 7, 2004			**November 2, 2004**		
Contract	Intrade Price	Predicted by Ranking Assumption	Predicted by Independ. Assumption	Intrade Price	Predicted by Ranking Assumption	Predicted by Independ. Assumption
Bush Greatplains	35.0	40.0	13.9	23.0	24.0	9.7
Bush OH+FL	56.9	60.5	38.1	37.0	51.1	27.5
Bush South	55.0	60.5	18.9	53.0	53.9	32.3
Bush Southwest	36.0	43.0	18.7	53.8	56.5	32.7
Kerry New England	53.7	58.0	33.5	70.0	69.0	57.1
Kerry Rustbelt	32.0	37.0	14.0	42.5	48.9	30.9
Kerry Westcoast	63.5	63.7	41.5	87.5	90.0	84.4

Notes:
- Greatplains: IA, KS, MN, NE, ND, OK, SD, & TX.
- South: SC, MS, FL, AL, GA, LA, TX, VA, AR, NC, & TN.
- Southwest: NV, NM, UT, & CO.
- New England: CT, RI, ME, VT, MA, & NH.
- Rustbelt: PA, OH, & MI.
- Westcoast: CA, OR, & WA.

the independence assumption is simply the product of the state prices (after dividing each price by 100 and multiplying the final product by 100).

It is clear from Table 4 that the predictions are much closer under the ranking assumption than under the independence assumption. The worst case for the independence assumption is Bush South, where for September 7 the ranking- assumption price is 60.5, the price for Florida, and the independence-assumption price is 18.9. These compare to the actual price of the contract of 55.0. The only weak case for the ranking assumption is Bush OH+FL for November 2, where the contract price is 37.0 and the price predicted by the ranking assumption is 51.1. Although the results in Table 4 have to be taken with some caution because the markets are thinly traded, they are supportive of the view that the Intrade traders are using the ranking assumption to price the combination contracts.

Table 5
Intrade Data on the National Contract

	9/7	9/21	10/5	10/19	11/2
National Contract	60.2	70.0	60.0	58.5	55.5
Pivotal State	60.5	70.0	63.5	57.8	[a]51.1
	FL	FL	FL	OH	OH

[a]Bid/ask spread was 50.0/55.5.

Table 5 shows the price of the national contract on each of the five days and the price of the pivotal state. Remember that under the ranking assumption the two prices should be the same. The table shows that the prices are quite close. On the last day the prices differ by 4.4, but the bid/ask spread for Ohio was quite large, and so the Ohio price may not be reliable.

Turning now to the 2006 Senate election, a widely traded contract on Intrade was one that stated that the Republicans would retain control of the Senate. This would have happened had the Republicans taken one of the first six states in Table 3 plus Tennessee (which according to Intrade was not close). If traders were using the ranking assumption, the price of this contract should have been 46.5, one minus the price for Virginia (using the last price). Under the assumption that the probabilities of the first six states are independent (and everything else certain), the price of the contract should have been 88.6, one minus the product of the six probabilities (using the last prices). The actual price at 6:00 AM was 66.4 using the last price and 67.2 using the average of the bid and ask prices (bid was 66.4, ask was 67.9). So the market price was almost exactly halfway between the ranking assumption price and the independence assumption price. The actual price is in fact consistent with the use of independence assumption for Missouri and Virginia and certainty otherwise. In this case the price should have been 69.6, one minus the product of the two probabilities (using the last prices), which is close to the actual. So in this case the traders were not using the ranking assumption to price the Republican-control contract. They, of course, should have, given that the actual results were exactly as the joint hypothesis of a correct Intrade ranking and the ranking assumption predicted. Those who ignored the ranking assumption and bought the contract (assuming, say, independence for Missouri and Virginia) lost.

5 Political Party Responses to Uncertainty

Estimation Errors

This section shows that if the ranking assumption holds, the two political parties in a presidential election should spend money in only a few states. It is first necessary to consider what it means within the context of this paper for the prices in Table 1 to change across time and in some cases to change the ranking of the states. It is important to realize that these changes, even changes in ranking, are not inconsistent with the ranking assumption because the assumption pertains only to the ranking on the last day.

Let p_i denote the probability that Bush wins state i on election day, which is the percent of the n conditions of nature in which Bush wins state i. Assume that these probabilities are estimated precisely by the Intrade prices on the day before the election—the prices in the last price column in Table 1.

Consider the prices on September 7, about two months before the election. Let

$\hat{p}_{it}$ denote the price for state i on date t, where in this case t is September 7. Let u_{it} denote the estimation error for state i and date t:

$$u_{it} = \hat{p}_{it} - p_i. \tag{1}$$

For t equal to September 7, u_{it} for a given state is the difference between the first price column in Table 1 and the last price column. Surprises that happen between, say, September 7 and election day will change the estimated probabilities (and thus prices) as people update their views about the conditions of nature that will exist on election day. A surprise negative performance by Bush in the debates would likely lower nearly all the estimated probabilities. If all the estimated probabilities fell by the same amount, there would be no change in the ranking. The fact that the ranking in Table 1 changes somewhat over time means that some surprises are state specific. There are thus state specific components in u_{it} in (1).

Stochastic Simulation

Before considering the spending strategy of the two parties, it will be useful to examine the effects of state-specific variation in the estimation errors. This is done in Table 6 using stochastic simulation. To focus on state-specific variation, the errors are taken for the simulation work to be uncorrelated across states. The states used are the 13 states with prices between 30.0 and 70.0 on September 7. For the results in Table 6 t is September 7. For each state i, u_{it} is assumed to be normally distributed with mean 0 and variance σ^2. σ^2 is assumed to be the same across states.

Table 6
Stochastic Simulation Results
Data for September 7, 2004

	Value of σ					
	0.00	0.01	0.02	0.03	0.04	0.05
$p_v^{(k)}$						
median	.600	.597	.592	.588	.582	.576
minimum	.600	.559	.522	.481	.439	.409
.05	.600	.583	.567	.551	.533	.515
# times pivotal state						
WV	0	0	0	48	111	180
MO	0	0	9	91	254	400
OH	0	50	705	1416	1962	2199
FL	0	3560	4185	4278	4185	4057
NV	10000	6218	4113	2913	2236	1814
WI	0	172	988	1254	1235	1197
PA	0	0	0	0	8	84
IA	0	0	0	0	3	26
NM	0	0	0	0	2	18
NH	0	0	0	0	4	10
MN	0	0	0	0	0	13
OR	0	0	0	0	0	1
MI	0	0	0	0	0	1
# times pivotal state or above						
WV	10000	10000	9999	9978	9906	9783
MO	10000	10000	10000	9982	9927	9807
OH	10000	10000	10000	10000	9997	9955
FL	10000	10000	10000	10000	9996	9943
NV	10000	9819	8733	8122	7869	7753
WI	0	248	2100	3616	4556	5208
PA	0	0	0	0	10	101
IA	0	0	0	0	12	104
NM	0	0	0	0	11	97
NH	0	0	0	0	8	50
MN	0	0	0	0	1	21
OR	0	0	0	0	0	2
MI	0	0	0	0	0	1

- The prices (base probabilities) from Table 1 for September 7 are: WV 67.7, MO 67.0, OH 63.0, FL 60.5, NV 60.0, WI 57.0, PA 43.4, IA 43.0, NM 43.0, NH 42.0, MN 40.0, OR 36.3, MI 33.0.
- 10000 trials per value of σ.
- $p_v^{(k)}$ = probability of winning the election for the kth trial, which is the probability of winning the pivotal state.
- .05 for $p_v^{(k)}$ means the value below which 5 percent of the trial values lie.

The stochastic-simulation experiments were performed as follows. For each trial 13 errors were drawn from the $N(0, \sigma^2)$ distribution, one per state, where σ varied from zero for the first experiment to 0.05 for the sixth experiment. Consider a given experiment, i.e., a given value of σ. Let $u_{it}^{(k)}$ denote the error drawn for state i on the kth trial. The probability for state i on the kth trial was computed as:

$$p_{it}^{(k)} = \hat{p}_{it} + u_{it}^{(k)}. \tag{2}$$

In this context $\hat{p}_{it}$ is the "base" probability. For each trial k the values of $p_{it}^{(k)}$ were ranked, the pivotal state was determined,[6] and its probability, denoted $p_{pt}^{(k)}$, was recorded. This was done 10,000 times, resulting in 10,000 values of $p_{pt}^{(k)}$. The number of times a particular state was the pivotal state was also recorded, as was the number of times a state was above the pivotal state. Presented in Table 6 are the minimum value of $p_{pt}^{(k)}$, the value below which 5 percent of the trial values lie, and the median. Also presented are the number of times each state was pivotal and the number of times each state was pivotal or above the pivotal.[7]

The results in Table 6 are easy to explain. When the variance is zero, Nevada is always pivotal and the probability of winning the election is always .600.[8] As the variance increases, more and more states are sometimes pivotal or above the

[6]For this work 270, not 269, was taken to be the number of electoral votes needed to win.

[7]It can be the case in the stochastic simulations that $p_{it}^{(k)}$ for a particular state i is greater than the base probability for states above the highest ranked state used (West Virginia) or less than the base probability for states below the lowest ranked state used (Michigan). This does not matter for the results, however, because the solutions that matter are around the pivotal state. The stochastic simulation could have been set up using all the states, but, as just noted, this is not necessary. If all states were used, the assumption that the variance of the error is the same across states would have to be changed. The variance is obviously smaller when the base probability is near one or zero than when it is near one half.

[8]In Table 4 Florida is listed as the pivotal state for September 7, whereas in Table 6 Nevada is listed as pivotal. This difference is due to the use of 270 electoral votes to win rather than 269.

pivotal. The median of $p_{pt}^{(k)}$ falls from .600 when σ is zero to .576 when σ is 0.05. The median falls because, except for Wisconsin, the states below Nevada have base probabilities that are considerably below .600. There is not symmetry around .600, and so negative draws for states above Nevada are on average not completely offset by positive draws for states below Nevada. When the calculations were repeated using .570 for the base probabilities for the states below Wisconsin (instead of the values in Table 1 for September 7), the median of $p_{pt}^{(k)}$ rose as the variance increased. For $\sigma = 0.01$ the median was .597. The values of the median for the increasing values of σ were, respectively, .598, .600, .603, and .605.

When σ is zero, i.e., no state-specific variation, all that matters in terms of predicting the probability of winning the election is the probability for the pivotal state. It does not matter, for example, how much larger the probabilities for the states above the pivotal state are or how much smaller the probabilities for the states below the pivotal state are. As just seen, this changes when σ is non zero—the sizes of the probabilities around the pivotal state now matter.

The stochastic simulations were repeated using the September 21 data (t = September 21), and the results are presented in Table 7. These results are similar to those in Table 7, although with higher probabilities, except that some states are now never pivotal nor above the pivotal. The fact that the base probabilities for Iowa and New Hampshire have risen substantially leads to these states doing all the extra work. Even with its 21 electoral votes, Pennsylvania is never used.

Table 7
Stochastic Simulation Results
Data for September 21, 2004

	Value of σ					
	0.00	0.01	0.02	0.03	0.04	0.05
$p_v^{(k)}$						
median	.699	.694	.688	.680	.673	.667
minimum	.699	.658	.617	.576	.534	.492
.05	.699	.680	.660	.642	.623	.606
# times pivotal state						
MO	0	0	0	0	2	7
WV	0	0	4	78	187	296
OH	0	219	1100	1733	2103	2333
FL	0	4553	4264	4016	3870	3819
NV	10000	5228	4610	3898	3265	2648
WI	0	0	22	268	532	743
IA	0	0	0	3	27	80
NH	0	0	0	4	14	74
# times pivotal state or above						
WV	10000	10000	10000	9998	9980	9908
MO	10000	10000	10000	10000	10000	10000
OH	10000	10000	10000	10000	9999	9971
FL	10000	10000	10000	10000	10000	10000
NV	10000	10000	9977	9683	9277	8838
WI	0	0	48	624	1543	2456
IA	0	0	0	5	74	285
NH	0	0	0	7	54	300

• See notes to Table 4.
• The prices (base probabilities) from Table 1 for September 21 are: MO 85.0, WV 77.0, OH 72.0, FL 70.0, NV 69.9, WI 62.0, IA 55.0, NH 55.0, PA 43.0, MN 40.5, NM 40.0, OR 35.0, MI 29.9.
• PA, NM, MN, OR, and MI were never used.

Campaign Spending

The insights from Tables 6 and 7 can now be used to examine campaign spending across states. Each possible condition of nature on election day is based on everything that has happened up to the day of the election. "Everything" includes all the campaigning that has been done in each state. After all the campaigning is over, the ranking assumption says that there is no possible condition of nature in which Bush wins a state ranked below a state he loses. This is not to say, of course, that campaigning has no effect on the possible conditions of nature. It is just that once campaigning is over, the ranking assumption holds.

Consider now the strategy of the Republican party on some date t before the election. Assume for now that the Republican party does not take into account any Democratic-party response to its actions. $\hat{p}_{it}$ is the market's estimate at date t of what the actual probability will be on election day (p_i). This estimate obviously takes into account market participants' views about how much campaigning there will be in each state. Let r^e_{it} denote the market's expectation at date t of the amount the Republican party will spend in state i between date t and election day, and let d^e_{it} the similar variable for the Democratic party. The following equation is then postulated:

$$p_i = \hat{p}_{it} + f_i(r_{it} - r^e_{it}) - g_i(d_{it} - d^e_{it}) + u_{it} \tag{3}$$

where r_{it} is the actual amount the Republican party spends in state i between date t and election day and d_{it} is the similar variable for the Democratic party. Equation (3) says that spending in a state affects the probability of winning the state. The Republican party faces a budget constraint that the sum of r_{it} across all the states cannot exceed some amount, and similarly for the Democratic party.

Assume that decisions are being made on date t equal to September 7, so t is fixed, and assume for now that d_{it} does not respond to changes in r_{it}. If the Republican party wants to maximize the probability of winning the election, what should it do? Consider first the case in which the variance of u_{it} in equation (3) is zero for all i. In this case under the ranking assumption the Republican party simply maximizes the probability of winning the pivotal state. In Table 1 for September 7 the pivotal state is Nevada (assuming 270 electoral votes needed to win), which has a price of 60.0. The state above it is Florida, with a price of 60.5. The next state is Ohio, with a price of 63.0, and the next state is Missouri with a price of 67.0. To take an example, say the Republican party's budget constraint is such that the party can spend in Nevada, Florida, and Ohio to raise p_i to 65.0 each. The probability of winning has thus increased from .60 to .65, and there has been spending in just three states. (In this example there would be in the end no conditions of nature on election day in which Bush won one or two of these states and lost the other.)

Consider next the case in which the variance of u_{it} is not zero. Remember that the u_{it} are state-specific errors of estimation. On date t (September 7) the Republican party knows that it can change the actual probabilities that will exist on election day, but when there are estimation errors it does not know the actual values that will exit. What should be the objective of the party in this case? Go back to the stochastic-simulation results in Table 6 and assume that the 13 states in the table are in play. Let r_t denote the vector of the 13 r_{it} values, and let u_t denote the vector of the 13 u_{it} values. Given r_t and u_t, it is straightforward to compute the probability that the Republican party wins the election. The values of p_i can be computed from equation (3) (assuming also knowledge of r_{it}^e, d_{it}, and d_{it}^e) and then the values ranked to determine the pivotal-state value. For the given value of r_t this can be done, say, for 10,000 draws of u_t. This gives 10,000 values of the probability of winning the election, from which summary measures like those in

Table 6 can be computed.

One can think of the Republican party considering many values of r_t and for each value computing 10,000 probabilities and summary measures like those in Table 6. Its objective might be to choose r_t to maximize the median of the probability values, the minimum of the values, or the value below which 5 percent of the trial values lie. This last option means that there would be a 95 percent chance that the actual probability of winning on election day is above the maximized value. Whatever is maximized, Table 6 shows that when the variance of the errors is zero the optimal strategy for the party would be to allocate some of its spending to states below Nevada, the pivotal state when the variance of the errors is zero. Some states that are below Nevada now have, depending on the draw for u_t, some chance of being pivotal, and so it would be optimal to spend something on these states.

The addition of uncertainty has thus increased the number of states in which spending is done. Table 6 shows that as the variance of the errors increases, the number of states that are sometimes pivotal increases. Thus, the larger the variance, the larger the number of states in which spending is done. It is still the case, of course, that in most states no spending is done.

Consider finally the Democratic-party response to a Republican-party move, i.e., relax the assumption that d_{it} is fixed. . In any given presidential election the two parties generally have similar resources and similar information. It also seems likely that the effects of spending on votes are similar between the two parties. If there is complete symmetry between the two parties and, say, the Republicans move first, then the Democrats can merely offset whatever the Republicans do. In practice this seems to be roughly the case. Both parties focus their spending on the swing states and come close to matching each other by state in terms of number of visits by the candidates and advertising spending. If one party begins to do more in a key state, the other party tends to respond. Also, there is essentially no spending in many states, which, as discussed next, is consistent with the ranking assumption but not the independence assumption.

No attempt is made in this paper to set up a formal game between the two parties under the ranking assumption. This is a possibly interesting area for future work. With a probability structure like that in Table 1, where many states are close to zero or 100, it seems clear from the results in Table 6 that if a game is set up using the ranking assumption, there are likely to be many states in which there is no spending by either party. This is contrary to results in the literature that are based on the independence assumption. In the model of Snyder (1989), for example, spending is high in states that are close and that have a high probability of being pivotal, but there is some spending in all states. The same is true for the model in Strömberg (2002). In the model of Brams and Davis (1974) there is spending in all states, where spending is in proportion to the 3/2's power of the number of electoral votes in each state.

6 Conclusion

This paper provides an interpretation of the uncertainty that exists at the beginning of the day of an election as to who will win. It is based on the theory that there are a number of possible conditions of nature that can exist on election day, of which one is drawn. Political betting markets like Intrade provide a way of trying to estimate this uncertainty. Polling standard errors, on the other hand, do not provide estimates of this type of uncertainty. They estimate sample-size uncertainty, which can be driven close to zero with a large enough sample.

This paper also introduces a ranking assumption, which puts restrictions on the possible conditions of nature than can exist on election day. The joint hypothesis that the last-day Intrade ranking is correct and the ranking assumption is correct predicts the exact outcomes of the 2004 presidential election and the 2006 Senate election. Although not a test of the ranking assumption, there is evidence that the Intrade traders used the ranking assumption to price contracts in the 2004 presidential election. This was not the case, however, in the 2006 Senate election.

Under the assumption that the ranking assumption is correct, the stochastic simulation results in Section 5 show that the two political parties should spend only in a few states. The larger the variance of the estimation errors, the larger is the number of states in play, although even for large variances the number of states in play is small.

Finally, given the success of the ranking assumption in 2004 and 2006, it will be interesting to see how it does in the 2008 election. Regarding polling standard errors, this paper should not be interpreted as an attack on polls. All it says is that there is a type of uncertainty that is not estimated by polling standard errors but that can be estimated using political betting markets.

References

[1] Brams, Steven J., and Morton D. Davis, 1974, "The 3/2's Rule in Presidential Campaigning," *The American Political Science Review*, 68 (March), 113–134.

[2] Kaplan, Edward H., and Arnold Barnett, 2003, "A New Approach to Estimating the Probability of Winning the Presidency," *Operations Research*, 51 (January/February), 32–34.

[3] Leigh, Andrew, and Justin Wolfers, 2006, "Competing Approaches to Forecasting Elections: Economic Models, Opinion Polling and Prediction Markets," *The Economic Record*, 82 (September), 325–340.

[4] Manski, Charles F., 2004, "Interpreting the Predictions of Prediction Markets," February.

[5] Snyder, James M., 1989, "Election Goals and the Allocation of Campaign Resources," *Econometrica*, 57 (May), 637–660.

[6] Soumbatiants, Souren, Henry W. Chappell, Fr., and Eric Johnson, 2006, "Using State Polls to Forecast U.S. Presidential Election Outcomes," *Public Choice*, 127, 207–223.

[7] Strömberg, David, 2002, "Optimal Campaigning in Presidential Elections: The Probability of Being Florida," Seminar Paper No. 706, Institute for International Economic Studies, Stockholm University.

[8] Wolfers, Justin, and Eric Zitzewitz, 2004a, "Prediction Markets," *Journal of Economic Perspectives*, 18 (Spring), 107-126.

[9] Wolfers, Justin, and Eric Zitzewitz, 2004b, "Using Markets to Inform Policy: The Case of the Iraq War," June.