Prediction Market Accuracy in the Long Run*

Joyce Berg, Forrest Nelson and Thomas Rietz

Departments of Accounting, Economics and Finance
Henry B. Tippie College of Business
University of Iowa
Iowa City, Iowa 52242

Working Draft

June 2007

Abstract

"Prediction markets" are designed specifically to forecast events such as elections. Though such markets have been conduced for nearly twenty years, to date nearly all evidence on efficiency compares election eve forecasts to final pre-election polls and actual outcomes. Here, we present evidence that prediction markets outperform polls in the long run. For all available national polls, we ask whether the poll or a contemporaneous Iowa Electronic Markets vote share market prediction is closer to the eventual outcome for the two-major-party vote split. We compare market predictions to 964 polls over the five Presidential elections since 1988. The market is closer to the eventual outcome 74% of the time. Further, the market significantly outperforms the polls in every election when forecasting more than 100 days in advance.

I. Introduction

"Prediction markets" are designed and conducted for the primary purpose of aggregating information so that market prices forecast future events. These markets differ from typical, naturally occurring markets in their primary role as a forecasting tool instead of a resource allocation mechanism. For example, since 1988, faculty at the Henry B. Tippie College of Business at the University of Iowa have been running markets through the Iowa Electronic Markets (IEM) project that are designed to predict election outcomes. These represent the longest running set of prediction markets known to us. They have proven efficient in forecasting the evening and week before elections. Here, we show that, well in advance of the elections, these markets dominate polls in forecasting the election outcome.

Existing evidence (e.g., Berg, Forsythe, Nelson and Rietz, 2003, and references cited therein) shows excellent ex-post predictive accuracy for election prediction markets in the very short run (i.e., one-day-ahead forecasts using election eve prices). Extending a similar figure from Berg, Forsythe, Nelson and Rietz (2003) to include the 2004 election results, Figure 1 shows this accuracy on election eve. The average absolute percentage error in predicting the vote share for presidential candidates on election eve is 1.33 percentage points. Again, extending Berg, Forsythe, Nelson and Rietz (2003), our analysis shows that final poll errors average 1.62 percentage points.

Here, we present an analysis of the long-run forecasting ability of markets relative to polls. Results show that prediction markets are more accurate long-run forecasting tools than polls across elections and across long periods of time preceding elections (instead of just on election-eve). The basis for our statement is a simple one. We compare the market prediction for the two-party vote split to the poll prediction (normalizing to control for third party and undecided votes in the poll). We do the comparison while the poll is in the field to give a fair test. We simply ask how often the market prediction is closer to the eventual outcome than the

poll. Aggregating over 965 polls over the five Presidential elections since 1988, the market is closer to the eventual outcome 74% of the time. Further, the market significantly outperforms the polls in every election when forecasting more than 100 days in advance.

In the next section, we briefly describe prediction markets in general and the specific election markets we study. Then, we present our results and end with concluding remarks.

II. Prediction Markets

Since Hayek (1945), economists have recognized that markets have a dual role. They allocate resources and, through the process of price discovery, they aggregate information about the values of these resources. The information aggregation role of some markets seems particularly apparent. For example, corporations cite the value of their stock as the consensus judgment of their owners about the value of the corporation's activities. Increasingly, corporations reward managers based on this value measure. Futures and options markets aggregate information about the anticipated future values of stocks and commodities. If it is true that futures prices are the best predictors of actual future spot prices (as the "expectations hypothesis" asserts), then futures prices constitute forecasts.[1] For example, Krueger and Kuttner (1996) discuss how the Federal Funds futures contract can be used to predict future Federal Funds rates and, hence, future Federal Reserve target rates.

In most markets, if prediction uses arise, they do so as a secondary information aggregation role. However, some recent markets have been designed specifically to exploit their information aggregation characteristics for use as dynamic forecasting systems. Examples

[1]Debate over the ability of futures markets to forecast future prices extends back to Keynes (1930) and Hicks (1946). Many of the arguments result from the secondary nature of information aggregation in these markets. The early "normal backwardization" versus "contago effect" arguments were based on relative power of speculators and hedgers. Today, the idea that "risk neutral" probabilities used to price futures and options differ from the "true" underlying probabilities result from relative levels of hedging demand in the markets. While the IEM markets discussed below may be subject to price deviations due to hedging activities, the narrow scope of the IEM markets, the small size of investments and analysis of individual traders (e.g., Forsythe, Nelson, Neumann and Wright, 1992, and Forsythe, Rietz and Ross, 1999) all lead us to conclude that hedging activities do not affect IEM prices significantly.

of such “prediction markets” include numerous markets run under the Iowa Electronic Markets (designed to predict elections, other political events, movie box office receipts, corporate earnings, returns, stock prices, etc.; see Forsythe, Nelson, Neumann and Write, 1992 and Berg, Forsythe, Nelson and Rietz, 2003 for more detailed descriptions), similar markets run in other countries (usually designed to predict election outcomes) and markets cited in Plott (2000) (designed to predict sales at a large corporation). While the majority of such markets are run with cash payoffs, some similar Internet “games” have been run using fictitious currency with prize contests as motivation. These include the Foresight Exchange (http://www.ideosphere.com) with “payoffs” tied to a wide range of social, political and scientific events/issues, the Hollywood Stock Exchange (http://www.hsx.com) with “payoffs” tied to movie box office take and the (apparently now defunct) Major League Market (http://majorleaguemarket.com) with “payoffs’ tied to the performance of teams and athletes and similar markets with contract “de-listing values” (i.e., liquidating “payoffs”) tied directly to predictable events (in contrast to vague notions of “popularity” of contracts).

Prediction markets, such as the Iowa Electronic Markets (IEM), represent an important advance in forecasting. The idea is simple: trade contingent claims in a market where the claims pay off as a function of something one is interested in forecasting. If structured correctly, the prices should reflect the expected payoffs to the claims. This relationship can be used for forecasting. For example, the IEM’s vote-share markets trade contracts with payoffs that equal $1 times the relative percentages of the vote taken by candidates in an upcoming election. Prices should converge to the market’s expectation of relative vote shares. Though simple in concept, such markets act as complex, dynamic, interactive systems that incorporate information in new ways. Through the actions of traders, prediction markets aggregate information from individuals, incorporate polls and other sources of information and weight all of this information through the price formation process. They compete directly with, and potentially

use as information, traditional methods of forecasting such as polls, econometric modeling and marketing surveys.

Here we ask a simple question: Well in advance of the event, are prediction markets closer to eventual outcomes on average than the natural alternative? This extends the usual measure of predictive accuracy, which is based on election-eve market forecasts and final polls. To answer this question, we compare (1) the forecasts of IEM prediction markets run to predict vote shares of candidates in United States Presidential elections since 1988, (2) contemporaneous poll results and (3) the eventual outcomes of the elections. We simply ask which is closer to the eventual outcome: a poll or a price-based forecast taken from a market at the time the pollster was in the field. On average, we find that the markets are closer than polls to the eventual outcome.[2]

III. The Iowa Electronic Markets Presidential Vote Share Markets

The prediction markets we study in this paper are the Iowa Electronic Markets (IEM) Presidential Vote Share markets run since 1988. They are the longest running set of formal prediction markets known to us. The IEM prediction markets are computerized, electronic, real-time exchanges where traders buy and sell futures contacts with payoffs based on election outcomes. Traders entering the market are allowed to invest between $5.00 and $500.00. Because real money is used , traders are subject to the monetary risks and returns that result from their trading behavior.

Contracts in the vote share markets are designed to forecast the vote shares received by candidates. Contracts pay an amount equal to the fraction of the popular vote received by a

[2] This result may not be particularly surprising to pollsters since they do not claim to forecast the election, but instead simply measure current opinion. However, forecasting is a natural use of poll results and polls are the best alternative forecasts available.

candidate times $1.[3] Table 1 gives the specific contracts for the presidential vote-share markets run to date. Appropriate contract specification and normalization insured that the contract payoffs always summed to 1. Simple no-arbitrage arguments imply that market prices should reflect the traders' consensus forecast of the vote shares taken by each candidate[4]. Thus, vote-share markets provide point predictions about candidate vote shares. If, in advance of determining the election outcome, one thinks of each candidate's vote share as a random variable with some distribution, the vote-share market gives the mean of that candidate's vote-share distribution.

Table 2 shows statistics for the Presidential vote-share markets in the 1988 through 2004 elections.[5] The number of active traders in the vote share markets ranged from 155 in 1988 to 790 in the 2000 market. Overall volumes ranged from 15,826 contracts worth $8,123 in 1988 to 339,222 contracts worth $46,237 in 2004.

As a prediction system, the IEM differs from expert panels and polls in a number of respects. Instead of being a randomly selected, representative sample or a deliberately chosen panel, IEM traders are self-selected. People who are not interested either do not sign up or drop out. Further, the market does not equally weight traders' opinions in the price formation process. Instead, the market price is a weighted average which, through trading behavior and market dynamics, depends upon the traders' forecasts and the levels of confidence they have in

---

[3] In 1988, the contracts paid the vote share times $2.50.

[4] This can be shown in a variety of ways. See Malinvaud (1974) for the general equilibrium proof. One can also price the contracts as assets using CAPM and APT models. In each, $P_t = E(P_{t+h})/(1+k)^h$, where k (the required expected return) is the sum of the risk free rate and compensation for aggregate risk factors. Since the risk free rate is zero and there are no aggregate risk factors, the expected return on any given asset is zero. Alternatively, given that the expected market portfolio return is constrained to be zero by design, any factor risk premiums must be zero. Again, this makes for a zero expected return on any given asset. As a result, $P_t = E(P_{t+h})/(1+k)^h = E(P_{t+h})$. Even though traders cannot make the appropriate risk free hedges here (because they cannot trade the underlying fundamental asset), one might be tempted to use the modern portfolio theory futures pricing relationship: $F_{t+h} = E(P_{t+h})x(1+r_f)^t/(1+k)^t$, where F_{t+h} is the time t futures price for delivery at date $t+h$, $E(P_{t+h})$ is the expected future spot price of the underlying fundamental, r_f is the risk free rate and k is the required expected return determined by the risk of the futures position. Again, both the risk free rate and the required expected return are constrained to be zero. This gives: $F_{t+h} = E(P_{t+h})$.

[5] From 1992 on traders could participate in both vote-share and winner-takes-all markets. Here, we discuss the vote share markets only.

their forecasts as well as an untold number of factors like aggressiveness, risk aversion, timing, wealth, etc. Unlike polls or expert panels in which participants are asked for their independent opinions, each trader in the market sees the net effect of the beliefs of all other traders, and the time series of changes in those beliefs, and can alter his own perceptions accordingly. This makes the market more than a static, one-time prediction but rather a dynamic system that can respond instantaneously to the arrival of new information. Unlike polls that ask each respondent how he or she would vote if the election were held today, the market asks traders to forecast how everyone will vote in the actual upcoming election.

As an example of these differences, consider the demographics of IEM traders. A good poll would strive to collect responses from a random, representative sample. In contrast, IEM traders are self-selected and differ greatly from a representative sample of voters. In 1988, traders included only interested members of the University of Iowa academic community. In the other elections, traders included interested individuals from around the world. For example, in the 2000 vote-share market, 20% of the traders were from Iowa while Iowa only accounted for 1% of the nation's population in 2000. Men constituted 75% of the active traders but only 49% of the overall population (and slightly less of the voting population). Our traders are typically young, white, well educated and have high family incomes. Thus, IEM predictive accuracy relies heavily on a sample (in practice, a non-representative sample) of interested traders forecasting the behavior of the voting population at large. It does not depend on the traders themselves constituting a representative sample.

The question here is whether the IEM outperforms polls as a predictive system well in advance of the election outcomes.

IV. Performance Versus Polls

The IEM has conducted markets on five US presidential elections. Table 1 summarizes these markets. In 1988, a vote-share market predicted the popular vote shares taken by Bush,

Dukakis, Jackson and rest-of-the-field. In 1992, the IEM vote-share market was split between two sub-markets. One sub-market predicted the vote split between the Democrat (Clinton) and the Republican (Dole). A second sub-market predicted the split between the two major parties and Perot. The 1996 vote-share market predicted the vote split between the Clinton as the Democratic nominee and Dole as the Republican nominee. In 2000, the vote-share market forecast the election for the Democratic, Reform and Republican nominees (Gore, Buchanan and Bush, respectively). Finally, in 2004, the vote-share market predicted the split between the Democratic and Republican candidates.[6] In this paper, we focus on the vote-share markets and the vote splits among the Democratic and Republican candidates only because these are the most directly comparable to polls. We judge the accuracy of these market forecasts by comparing them to the actual election outcomes.

Polls used for comparison with the market include all nation-wide poll reports we were able to find for each of the five elections. For the three elections prior to 2000, reports were culled directly from news reports. For elections in 2000 and 2004, poll results were collected from http://pollingreport.com/. Poll reports based on samples of "Likely Voters" were chosen when possible; reports using "Registered Voters" was the second choice; if neither of those were reported then we used reports based on samples of "All Adults". Tracking polls are being reported with greater frequency in more recent elections. Those polls use rolling samples with, typically, N/3 new subjects added each day to replace the N/3 oldest subjects in the sample. Such overlapping samples result in a lack of independence from one day to the next. To avoid this dependence, we retain for analysis only every third report of a tracking poll, working backward from the last report so as to include data as close to the election as possible. Often pollsters will, in the same poll, ask for the favorite candidate from a broad list and again for the favorite from a narrow list, typically just two in the latter question. In such cases we use only the

[6] In 2004, the IEM ran dual races between Bush (as the Republican Nominee) and a set of possible democratic nominees. We only report the results for the Bush/Kerry race here.

result on the question with the broadest list of candidates. Polls reports with imprecise starting and ending dates, and polls conducted prior to the start of the market were excluded from the analysis. The final sample of polls included 59 polls from 1988, 151 in 1992, 157 in 1996, 229 in 2000 and 368 in 2004.

Figure 2 contains graphs of the margin of victory for the Electoral College winner as predicted by the polls and the market for the 5 elections. Market predictions are generated from closing prices (the last trade price before midnight each day). Poll outcomes represented in the graphs appear on the last day that polling took place for that particular poll, which is typically a day earlier than the release of the poll. For all elections except 2004, letters distinguish polls, indicating either the polling organization or the agent that requested and published the poll. For both market prices and polls, the outcomes are plotted as the normalized two-party vote margin. Thus, for example, the market outcome in Figure 2 for 1996 is computed as

$s^{VS}_{Clinton-Dole,t} := \frac{p^{VS}_{Clinton,t} - p^{VS}_{Dole,t}}{p^{VS}_{Clinton,t} + p^{VS}_{Dole,t}}$, where *s* designates the normalized spread and *p*'s designate closing market prices for candidates at time *t*.[7] The poll outcome is computed as

$s^{Poll}_{Clinton-Dole,t} := \frac{r^{Poll}_{Clinton,t} - r^{Poll}_{Dole,t}}{r^{Poll}_{Clinton,t} + r^{Poll}_{Dole,t}}$, where *s* designates the normalized spread and r's designate the percentages of poll respondents for candidates at time *t*.[8] In all five graphs, vertical lines indicate significant mid-campaign events (party conventions and debates) and a horizontal line shows the actual election outcome.

Several things are obvious from the five graphs. First, the markets present a very different picture of the elections than the polls. What the polls are measuring as voter sentiment at any particular point in time frequently differs greatly from what the market predicts will actually

[7] Closing prices are the last trade price before midnight each day. If no trade occurs in a day, the previous day's closing prices are carried over. Normalization adjusts for the possibility that non-synchronous trades lead to predictions that do not sum to 1 by adjusting each observation proportionately.

[8] The effect of this normalization on poll results is to allocate non-responses across the two candidates in proportion to the share of respondents choosing those candidates.

occur in the election. The market prediction often stays well above or below all contemporaneous polls for extended periods of time. During these periods, the market is typically closer to the final outcome than contemporaneous polls. Second, all five graphs reveal a striking volatility in poll outcomes, both in absolute terms and in comparison to the market. Polls on the same day by different organizations or subsequent polls by the same organizations frequently differ dramatically, generating differences that fall outside the quoted margins of error. In each election, we observe the well-known poll phenomenon of "convention bounce" (the tendency for a party to rise in the polls during that party's convention and then fall). These strong effects do not appear in the markets. Third, the market appears to forecast the election outcomes more closely than polls months in advance.

For a formal comparison of the accuracy of predictions from market prices and polls, we first pair each poll with a set of market prices from the IEM vote-share markets. These market prices are the midnight prices (closing prices) from the last day that the poll was in the field.[9] Next we normalize both the polls and the market prices so that each set of values sums to 1. For polls, we divide the poll share for the Democratic candidate and the poll share for the Republican candidate by the sum of the polls shares for the two candidates. Similarly, to normalize market prices, we divide the market price for the Democratic candidate and the market price for the Republican candidate by the sum of the market prices for the two candidates. We then ask a simple question: which was closer to the actual election outcome, the market or the poll. We use binomial tests to calculate the statistical significance of our results.

Table 3 gives the results of our analysis. In each election except 1988, the market significantly out performed polls overall, in the last 100 days before the election, the last 65 days

[9] We do this to compare what would be predicted from market prices to what would be predicted from polls. The poll prediction could not be made until after all poll results were collected. So, we use a comparable market price – the market price from the last day the poll was in the field. Note that choice means traders in the market would not yet have had access to the results from that particular poll. IEM prices are recorded at midnight each day, so our prices are the midnight prices.

before the election and the last 31 days before the election, with magnitudes of the advantage of the market over polls ranging from 63% to 87%. The market also out performs polls during the last 5 days before each of the five elections, though the number of polls in each election is too small for the binomial tests to have much power. However, aggregating cross all five elections, the markets significantly outperform polls in the last week before the election, with an advantage to the market in 68% of the cases. Overall, the market prediction is closer than the poll response to the ultimate election outcome 74% of the time.

Table 4 breaks the observations into non-overlapping time periods. This highlights the markets' impressive performance well in advance of the elections. All of the markets, including the smallest market in 1988, significantly outperformed polls that were conducted more than 100 days before the election. If anything, these results suggest that the market improves in relative (not absolute) accuracy the longer the time until the election.

There are several important issues in deciding how best to compare the accuracy of market prices and polls. We address three in robustness checks: which market price to compare to a particular poll, which poll to compare to a particular market price, and how to compare results across years. In our tests above, we use the midnight price from the last night each poll is in the field. However, it generally takes several days to create a set of poll data so that meaningful poll data is available only at the end of the poll. But, market prices are formed continuously and available at any point in time. To see whether our selection of market price makes a difference, we also compare polls to two other market price measures: a market price that is the average closing market price the days the poll was in the field, and a market price that is a time-weighted average of prices resulting from trades that occurred while the poll was in the field.

In our tests reported above, each set of polling data is treated as an independent comparison of a poll to the market price. However, there are many polling organizations and often they are generating contemporaneous polls. For instance, our set of polls includes 10

distinct polls released on 10/31/04, an extreme case. In our analysis, we counted each of these polls as a unique observation, comparing it to the same market price. In essence, some market prices are "counted" more than once in making our accuracy comparisons. To see whether this makes a difference, we perform two additional tests: we compare each market price used in our original tests to the average of the polls whose last field days correspond to that price, and we use regression which clustering on days.

Our tests also weight each poll equally, regardless of the polling organization and election year. But, the number of polls and the number of polling organizations differs across election years. For instance, from 1988 to 2004, there was a 523% increase in the number of polls (59 to 368). To see whether this is driving our results, we perform two additional tests: we track the same polling organizations across election years (so that the firms in the samples are not changing) and we equally weight years rather than polls.

V. Concluding Remarks

Previous research has shown the absolute and relative accuracy of the markets at very short horizons (1 day to 1 week). The evidence here shows that the markets are also accurate months in advance and do a markedly better job than polls at these longer horizons.

References

Berg, JE, R Forsythe, FD Nelson and TA Rietz, 2003, "Results from a Dozen Years of Election Futures Markets Research," forthcoming in CR Plott and V Smith (Eds) Handbook of Experimental Economic Results, North Holland, Amsterdam.

Forsythe, R, FD Nelson, GR Neumann and J Wright, 1992, "Anatomy of an Experimental Political Stock Market," American Economic Review, 82:1142-1161.

Forsythe, R, TA Rietz and TW Ross, 1999, "Wishes, Expectations and Actions: Price Formation in Election Stock Markets," Journal of Economic Behavior and Organization, 39: 83-110.

Hayek, FA, 1945, "The Use of Knowledge in Society," The American Economic Review, 35, 519-530.

Hicks, JR, 1946, Value and Capital, 2nd ed., London: Oxford University Press.

Keynes, JM, 1930, Treatise on Money, 2nd ed., London: Macmillan.

Krueger, JT and KN Kuttner, 1996, "The Fed Funds Futures Rate as a Predictor of Federal Reserve Policy," Journal of Futures Markets, December 1996.

Plott, CR, 2000, "Markets as Information Gathering Tools," Southern Economic Journal, 67, 1-15.

Figures

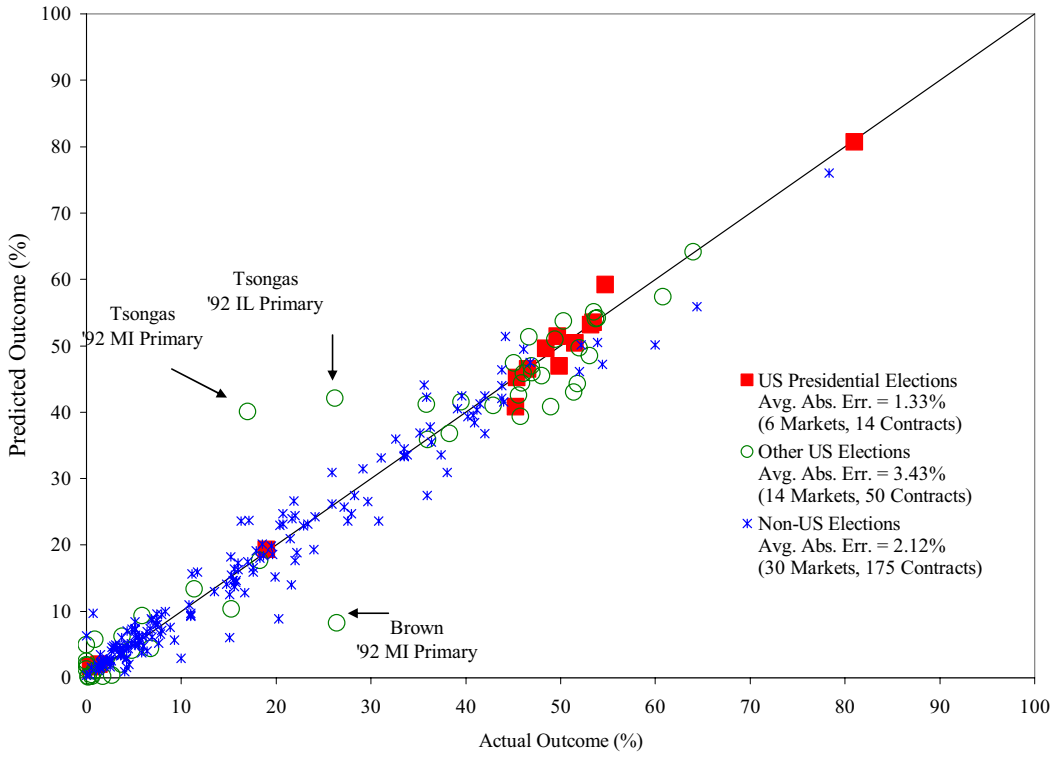

Figure 1: Election Eve Forecast Vote Shares and Actual Outcomes

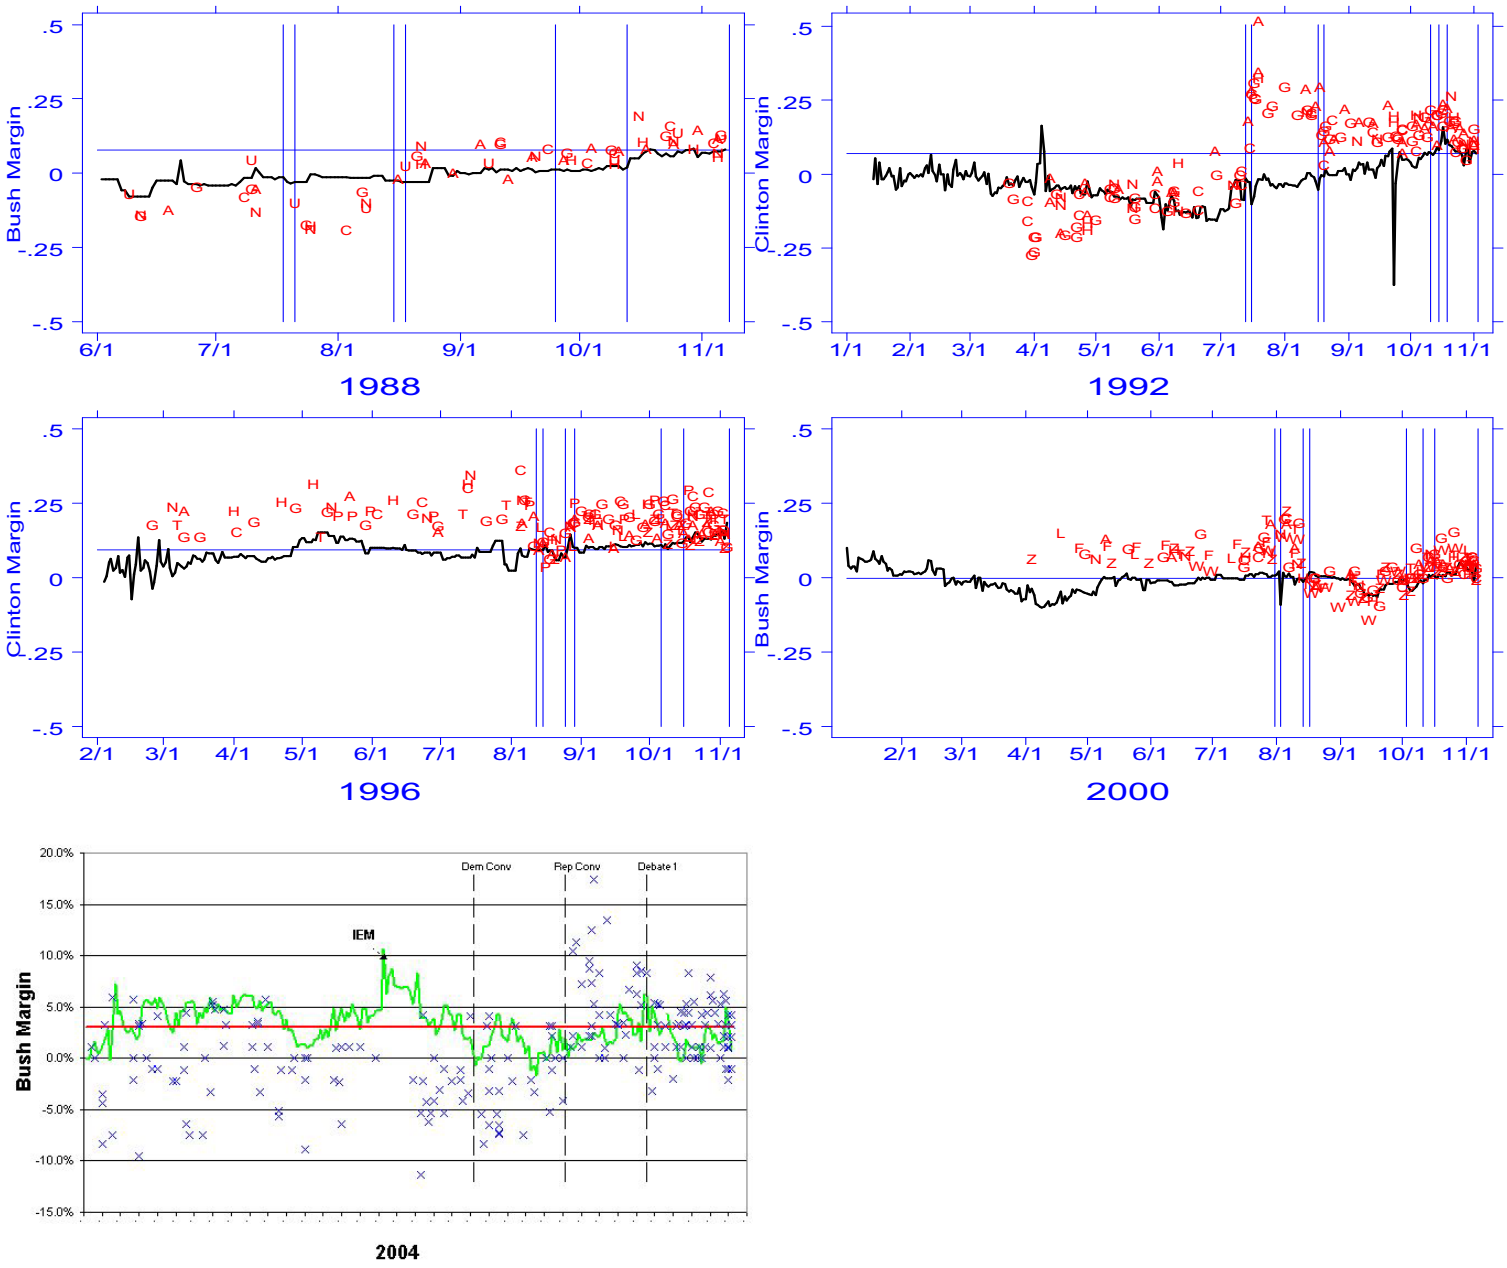

Figure 2: Predicted Vote Share Margins for Market and Polls

Verticle axis is vote margin; horizontal axis is date; the margin implied by the market is the solid moving line; poll margins are represented by letters (A=ABC, C=CBS, T=CNN, F=Fox, G=Gallup, H=Harris, L=Hotline, N=NBC, W=Newsweek, Z=Zogby); horizontal line at election outcome, vertical lines at beginning and end of conventions, debate days and election day.

Tables

Table 1: Contracts Traded in IEM
Presidential Vote-Share Election Markets

Year	Final Contracts Traded[1]	Number of Weeks Open Prior to Election
1988	Bush Dukakis Jackson Rest-of-field	23
1992	Democrat Republican Perot Democrat and Republican	43
1996	Democrat Republican	39
2000	Democrat Reform Republican	45
2004	Democrat Republican	37

[1]In 2004, contracts traded in a number of potential democratic nominees versus Bush. Here, we report only on the Bush/Kerry race.

Table 2: Summary of IEM Presidential Vote-Share Election Market Activity

	1988	1992[1]	1996[1]	2000[1]	2004[1]
Market Statistics					
Opening Date	6/1/1988	1/10/1992	2/4/1996	1/5/2000	2/21/2003
Election Date	11/8/1988	11/3/1992	11/5/1996	11/7/2000	11/2/2004
Weeks Open	23	43	39	44	45
Trader Investments	$4,976	$79,356	$200,000[2]	$210,633.00	$355,281.00
Overall Market Activity					
No. of Active Traders	155	592	264	790	777
Contract Volume	15,826	78,007	23,093	46,820	339,222
Dollar Volume	$8,123	$21,445	$3,628	$13,694	$46,237
Activity in Last Week					
No. of Active Traders	54	114	41	104	124
Contract Volume	962	1,389	592	4,192	4,947
Dollar Volume	$1,924	$569	$312	$609	$2,476

[1]Traders may have been active in multiple markets. Investments were fungible across markets.
[2]Estimated.

Table 3: Cumulative Binomial Tests for Relative Accuracy of the Market And Contemporaneous Poll Predictions

Poll predictions come from major polls taken during the election and are the normalized two-party vote shares. The market predictions are the normalized two-party vote share predictions on the last day each poll was in the field collecting data. The binomial variable takes the value 1 if the market prediction is closer the actual election outcome and 0 otherwise. Each p-value is the exact binomial probability of a number of 1s that large or larger, given that number of trails and a hypothesized probability of 0.50.) The number of observations is the number of polls in the sample period. If multiple polls are released on the same day, the same market price is compared to each poll.

Days included in sample	Item	1988	1992	1996	2000	2004	all years
All (from the beginning of the market	Number of polls	59	151	157	229	368	964
	poll "wins"	25	43	21	56	110	255
	market "wins"	34	108	136	173	258	709
	% market	58%	72%	87%	76%	70%	74%
	p-value (1sided)	0.149	0.000	0.000	0.000	0.000	0.000
Last 100 Days	Number of polls	45	82	124	180	173	604
	poll "wins"	24	23	18	54	44	163
	market "wins"	21	59	106	126	129	441
	% market	47%	72%	85%	70%	75%	73%
	p-value (1sided)	0.724	0.000	0.000	0.000	0.000	0.000
Last 65 Days	Number of polls	34	62	91	141	145	473
	poll "wins"	19	15	15	52	41	142
	market "wins"	15	47	76	89	104	331
	% market	44%	76%	84%	63%	72%	70%
	p-value (1sided)	0.804	0.000	0.000	0.001	0.000	0.000
Last 31 Days	Number of polls	21	40	58	84	97	300
	poll "wins"	7	7	13	26	33	86
	market "wins"	14	33	45	58	64	214
	% market	67%	83%	78%	69%	66%	71%
	p-value (1sided)	0.095	0.000	0.000	0.000	0.001	0.000
Last 5 Days	Number of polls	6	6	11	25	30	78
	poll "wins"	0	1	4	8	12	25
	market "wins"	6	5	7	17	18	53
	% market	100%	83%	64%	68%	60%	68%
	p-value (1sided)	0.016	0.109	0.274	0.054	0.181	0.001

Table 4: Cumulative Binomial Tests for Relative Accuracy of the Market And Contemporaneous Poll Predictions for Specific Time Ranges

Poll predictions come from major polls taken during the election and are the normalized two-party vote shares. The market predictions are the normalized two-party vote share predictions on the last day each poll was in the field collecting data. The binomial variable takes the value 1 if the market prediction is closer the actual election outcome and 0 otherwise. Each p-value is the exact binomial probability of a number of 1s that large or larger, given that number of trails and a hypothesized probability of 0.50.) The number of observations is the number of polls in the sample period. If multiple polls are released on the same day, the same market price is compared to each poll.

Days included in sample	Item	1988	1992	1996	2000	2004	All years
More than 100 Days	Number of polls	14	69	33	49	195	360
	poll "wins"	1	20	3	2	66	92
	market "wins"	13	49	30	47	129	268
	% market	93%	71%	91%	96%	66%	74%
	p-value (1sided)	0.001	0.000	0.000	0.000	0.000	0.000
66-100 Days	Number of polls	11	20	33	39	28	131
	poll "wins"	5	8	3	2	3	21
	market "wins"	6	12	30	37	25	110
	% market	55%	60%	91%	95%	89%	84%
	p-value (1sided)	0.500	0.252	0.000	0.000	0.000	0.000
32-65 Days	Number of polls	13	22	33	57	48	173
	poll "wins"	12	8	2	26	8	56
	market "wins"	1	14	31	31	40	117
	% market	8%	64%	94%	54%	83%	68%
	p-value (1sided)	1.000	0.143	0.000	0.298	0.000	0.000
6-31 Days	Number of polls	15	34	47	59	67	222
	poll "wins"	7	6	9	18	21	61
	market "wins"	8	28	38	41	46	161
	% market	53%	82%	81%	69%	69%	73%
	p-value (1sided)	0.500	0.000	0.000	0.002	0.002	0.000
Last 5 Days	Number of polls	6	6	11	25	30	78
	poll "wins"	0	1	4	8	12	25
	market "wins"	6	5	7	17	18	53
	% market	100%	83%	64%	68%	60%	68%
	p-value (1sided)	0.016	0.109	0.274	0.054	0.181	0.001

x