

Results from a Dozen Years of Election Futures Markets Research*

by

Joyce Berg**

Robert Forsythe***

Forrest Nelson***

Thomas Rietz****

March 2003

Working Draft for eventual publication in

The Handbook of Experimental Economics Results

CR Plott and VL Smith, Editors

Keywords: Iowa Electronic Markets, Prediction Markets, Election Markets, Political Stock Markets Forecasting, Polls, Elections.

JEL classifications: C19, C49, C8, C93, D4, E17, G13, H10

* For many helpful comments and discussions during the course of research summarized here, we thank Tom Gruca, George Neumann, Kent Oliven, Tom Ross and Jack Wright.

** Department of Accounting, Henry B. Tippie College of Business, University of Iowa, Iowa City, Iowa 52242.

*** Department of Economics, Henry B. Tippie College of Business, University of Iowa, Iowa City, Iowa 52242.

**** Department of Finance, Henry B. Tippie College of Business, University of Iowa, Iowa City, Iowa 52242.

Results from a Dozen Years of Election Futures Markets Research

1 Introduction and description of election futures markets

The Iowa Electronic Markets are small-scale,¹ real-money futures markets conducted by the University of Iowa Henry B. Tippie College of Business. In this review, we focus on the best known of these markets, The Iowa Political Markets.² Contracts in these markets are designed so that prices should predict election outcomes. The data set contains the results of 49 markets covering 41 elections in 13 countries.

The Iowa Electronic Markets operate 24-hours a day, using a continuous, double-auction trading mechanism. Traders invest their own funds, make their own trades, and conduct their own information search. The markets occupy a niche between the stylized, tightly controlled markets conducted in the laboratory and the information-rich environments of naturally occurring markets. By virtue of this design, the Iowa Markets provide data to researchers that is not otherwise available.

In addition to examining the accuracy of prices in these markets, we also compare the results of the national elections to a natural benchmark, polls, when available. Relative to polls, the markets rely on very different mechanisms for data collection and aggregation. Polls ask the question, “If the election were being held today, would you vote for the Democratic candidate or for the Republican candidate?” They rely on a representative sample of likely voters, truthful responses to the poll questions and classical statistics to arrive at their prediction of election outcomes. In the Iowa Markets, traders receive an explicit financial reward tied to correctly answering the question, “Who will everyone vote for on election day?” Traders are not a representative sample of likely voters; they are overwhelmingly male, well-educated, high income, and young (the average age is close to 30). In fact, we do not require our traders to be eligible to vote in the election.

¹ Investments are typically limited to a \$500 maximum per trader and generally average less than a tenth of this. Vote share markets have ranged in size from a dozen or so active traders to more than 500. Dollar and quantity volumes range up to \$21,445 and 78,007 contracts in the 1992 U.S. Presidential Vote Share market (see Berg, Nelson and Rietz, 2001).

² The Iowa Electronic Markets offers vote-share, seat-share and winner-takes-all election markets and markets based on other political outcomes, economic indicators and corporate earnings and returns. Here we focus on vote-share and seat-share election markets because, unlike other IEM offerings, they have both an accuracy criterion (election outcome) and recognizable performance benchmarks (polls).

2 Market Mechanism³

Each market is related to a specific future event, for instance a presidential election, and contains a set of contracts with liquidation values pegged to the outcome of the future event. Contracts enter into circulation by the voluntary purchase from the system of bundles of contracts that we call “unit portfolios,” or they are removed from circulation by sales of unit portfolios back to the system. These unit portfolios consist of one of each contract available in the market, and they are purchased from and sold to the exchange system for a fixed fee, which is the predetermined aggregate payoff to that portfolio. This use of unit portfolios ensures that the market operates as a zero-sum game and it permits the supply of contracts to be determined endogenously by the net number of unit portfolios that have been purchased by traders. Unit portfolios are employed only to place contracts in circulation; transactions among traders occur with individual contracts at prices determined by the participants.

Traders in these markets can place market orders (requests for immediate execution of a trade at current market prices) and limit orders (offers to buy (bid) or sell (ask) specified quantities at specified prices within some specified period of time). Limit orders are kept in queues ordered by price and time. Traders can withdraw their own outstanding limit orders any time before they trade or expire. The market information set available to traders consists of current best bid and ask prices and the last trade price. Traders do not know the quantity available at the best bid and ask. Nor do they know other entries in the bid and ask queues, except for their own bids and asks. This information appears on the trading screen as depicted in Figure 1. In addition, traders can access historical daily price information consisting of the quantity and dollar volume and the high, low, average and last trade prices.⁴

3 Results from share markets

In vote-share markets, the relative vote shares that candidates receive determine contract liquidation values. Typically, a particular contract will have a liquidating payoff equal to \$1 times the vote share received by the associated candidate. We insure that vote shares sum to 100% by either (1) trading a single contract associated with all minor-party vote shares (a “rest-of-the-field” contract) or (2) calculating vote shares based on fractions of the major party vote (e.g., the Democratic vote divided by the summed Democratic and Republican votes).⁵ Simple

³ Here, we highlight the important features of the markets. For more detailed descriptions, see Berg, Forsythe and Rietz (1996 and 1997) and Forsythe, Nelson, Neumann and Wright (1992).

⁴ While our markets are continuous, we report price information by 24-hour daily periods.

⁵ A prospectus details the particular method used in each market.

arbitrage arguments imply that contracts should trade at the expected liquidation values.⁶ Thus, in these markets, prices should equal expected vote shares.

In seat-share markets, contracts liquidate at values determined by the congressional or parliamentary seats allocated to parties in an election. Typically, a particular contract will have a liquidating payoff equal to \$1 times the seat share allocated to the associated party. We insure that seat shares sum to 100% by either (1) trading a single contract associated with all minor parties (a “rest-of-the-field” contract) or (2) calculating seat shares based on fractions of one party versus the all the other parties (e.g., the Republican share of seats versus seats held by non-Republicans).⁷ In these markets, prices should equal expected seat shares.

3.1 *Absolute market accuracy*

Efficiency evaluations are simple in vote share markets because they can be compared directly to the election outcome. To evaluate the ultimate predictive efficiency of the market, we compare the market predictions at midnight the evening before the election to the actual election outcome.⁸ Figure 2 extends a similar figure in Forsythe, Rietz and Ross (1999). It shows the absolute accuracy of 237 contract predictions in 49 markets run in 13 countries. In this figure, the bottom axis shows actual outcomes while the left axis shows market predictions. If all predictions lined up on the 45-degree line, the market would be perfectly accurate. Over-prediction errors lie above the line and under-prediction errors below.

Figure 2 shows no obvious biases in the market forecasts and, on average, considerable accuracy, especially for large, U.S. election markets.⁹ Berg, Forsythe and Rietz (1996) study how aggregate market characteristics affect accuracy in U.S. markets. Three factors explain most of the variance in accuracy. Presidential election markets perform better than (typically lower profile) congressional, state and local election markets. Markets with more

⁶ This arises because there is no aggregate risk in the markets. Thus, all agents can hold the well-diversified, “market” portfolio consisting purely of unit portfolios. Individual contracts can be priced from this portfolio and the risk/return tradeoff inherent in it. The return to holding unit portfolios is the same as the risk free rate (zero here). Thus, there is no risk premium in these markets. Because the risk premium is zero, the expected return for each risky asset must also be zero. This can only be true if all contracts are priced at their expected values. See Rietz (1998) for a more detailed explanation of how these results can be derived from general equilibrium arguments, from the capital asset pricing model and from arbitrage pricing theory.

⁷ Again, a prospectus details the particular method used in each market.

⁸ The clock on the market’s host computer determines “Midnight.” For markets run from the University of Iowa, this is U.S. Central Time. We choose midnight because expiring orders are removed from the queues just after midnight. This may cause substantial increases in the bid/ask spread and, in turn, the price volatility associated with a bid/ask bounce.

⁹ Biases might be predicted by arguments along the lines of Kahneman and Tversky (1979) or Jacobsen, Potters, Schram, van Winden and Wit (2000). The obvious outliers are associated with Tsongas’s and Brown’s showing in 1992 Primary elections

volume near the election perform better than those with less. Finally, markets with fewer contracts (i.e., fewer candidates or parties) predict better than those with more.

3.2 *Accuracy relative to polls*

Figure 3 shows how the market compares to polls for the subset of national elections in which we have poll data for comparison.¹⁰ Since market prices vary continuously, the question arises of which price to take as the prediction from the market. We include two measures: (1) the market price as of midnight on election eve and (2) the volume weighted average price of all transactions over the week before the election. The former incorporates all information available to traders as of that point in time but often reflects a great deal of volatility which results from the thinning out of queues on the last day of trading. The latter involves trades that are largely contemporaneous to the polls against which the market prices are being compared. Polling data in the figure represents the absolute errors of the final pre-election polls averaged across all candidates and across all of the major polling organizations for which data was available. The market outperformed polls in 9 of 15 cases according to each measure (election eve closing prices and last week average prices). Across all elections, the average poll error was 1.91% while the average market error was 1.49% and 1.58% by the two measures. In a few cases (the 1988 and 1992 U.S. Presidential elections) the market dramatically outperformed polls. The worst outcome, the 1996 U.S. Presidential election, is a peculiar one that gets additional attention below. In the majority of other cases, the market does about as well as the average poll, sometimes worse but often better, even if by a small margin.

Election eve outcomes are to some extent less interesting than predictions over the full course of the campaign. We notice a general tendency for the market to be both closer to eventual election outcomes and more stable than polls over the course of election campaigns. We illustrate this point in Figure 4 using the worst performing U.S. Presidential market as indicated by the election eve outcome relative to polls. This graph shows the relative stability of the market compared to polls over time, a feature typical of markets run to date. In the previous presidential elections, the market changed little or became slightly more accurate near the end of the election. However, in 1996, the market diverged from the correct outcome in the final days to close at midnight on the eve of the election with prices further from the election outcome than they had been since the Super Tuesday primaries in March. Apparently a large cash influx by new traders late in the campaign drove the price movement.

¹⁰ Typically, polls are designed to evaluate current opinions and not to predict future election outcomes. However, they serve as the obvious basis of comparison for the market predictions.

Nevertheless, over the majority of the time this market ran, its predictions were dramatically more accurate and stable than polls. This shows the value of the markets as longer run forecasting devices.

4 How and why do election futures markets “work?”

For the markets to work in theory, two features must be present. First, there must be enough traders so that the aggregate of their knowledge can forecast correctly the outcome of the election. Second, the market mechanism must facilitate aggregation of their disparate information so that the prevailing market price becomes a sufficient statistic for the traders’ collective information. Whether real markets can accomplish these two tasks is a behavioral question best answered by observing market dynamics and individual trader actions.

At the market level, Berg, Forsythe and Rietz (1996) show that larger, more active markets with fewer contracts are more accurate. Forsythe, Nelson, Neumann and Wright (1992), discuss the relationship between polls and market prices. They conclude that, while pre-election polls are obvious sources of information for market traders, market prices do not follow poll results. If anything, market prices predict changes in polls.

Forsythe, Nelson, Neumann and Wright (1992), Oliven and Rietz (2002) and Forsythe, Rietz and Ross (1999) all show that individual traders display considerable biases and often make mistakes in the largest and most accurate of the IEM markets. All three papers reconcile these observations with apparently high levels of market efficiency in a similar way. Specifically, the core group of traders that tend to set market prices appears less biased and error prone than typical traders. Forsythe, Nelson, Neumann and Wright (1992) show that typical traders have unrealistically optimistic beliefs for their preferred candidates (see Granberg and Brent, 1983, for a historical discussion of this bias). However, “marginal traders” (those who regularly trade or place bids and asks near the top of the queues) are much less prone to this bias. Oliven and Rietz (2002) show that typical traders often trade at a price that is not the most advantageous price for the trader or that violates arbitrage restrictions. Such “mistakes” are “irrational” because they decrease the trader’s payoff regardless of expectations or outcomes. However, “market makers” (those who actually set market prices by placing the best bids and asks) make mistakes much less often. Forsythe, Rietz and Ross (1999) study these issues further, relating the biases to psychological foundations. These results are what distinguish much of economics from the other social sciences. Marginal traders, not average traders, drive market prices and, therefore, predictions.

We also believe that the differences between election markets and polls give the markets an edge in prediction. Not only are the traders paid for correct guesses about the eventual election outcomes, but the market

information set also includes previous market outcomes, poll results and any other information deemed relevant by traders.

5 Other issues studied and future research potential

In addition to share markets, the IEM runs political “winner-takes-all” markets. Contracts in “winner-takes-all” markets liquidate at \$1 if an associated event occurs. For example, if a particular candidate wins an election, the associated winner-takes-all contract will pay \$1. Because of how we specify contracts and the event space, a full portfolio of all contracts in a market will always liquidate at \$1. Again, prices should equal expected values. Because of the \$0/\$1 payoff structure, prices should equal the probability of the \$1 liquidation. Thus, prices can be used to evaluate a candidate’s chances of winning an election, assess the relative viability of candidates and measure the impact of specific events on elections. In addition to winner-takes-all election markets, we have run such markets on various other political events. These include markets designed to predict: who party nominees will be in elections, who will control the houses of congress, whether particular bills or treaties will pass (e.g., NAFTA), whether countries will join the EMU, etc. These markets respond quickly to some events, but often appear unaffected by events that one might otherwise think should affect a campaign. This allows us to separate “surprises” from “news” that is anticipated. (For examples of using the markets to evaluate news events, see Forsythe, Nelson, Neumann and Wright, 1991 and 1992).¹¹

In addition to predicting outcomes and evaluating accuracy relative to polls, the IEM has been used to study a variety of other market related research topics. Forsythe, Nelson, Neumann and Wright (1992) and Oliven and Rietz (2002) study relationships between individual trader characteristics and actions. Oliven and Rietz (2002) also study arbitrage restrictions and violations. Forsythe, Rietz and Ross (1999) study price formation and psychological biases. Bondarenko and Bossarts (2000) study price dynamics and Bayesian updating. Beckmann and Werding (1996) compare call and continuous markets. Slemrod and Greimel (1999) study the relationship between IEM forecasts and bond markets. Plott (2000) and Ortner (1997 and 1998) uses similar markets to predict corporate events. Berg and Rietz (2003) show how such markets can serve for decision support.

¹¹Since 1993, the IEM has conducted markets based on economic indicators and corporate earnings and returns. These markets are limited to participants from academic communities and, along with the political markets, are used in many classes. Examples of the “earnings and returns” include: monthly contracts based on futures market returns in the computer industry and Microsoft’s stock price level, periodic markets designed to predict earnings announcements of companies and periodic markets designed to predict box-office earnings for movies. For an example of research using earnings and returns markets, see Bondarenko and Bossarts (2000).

Thus, the IEM and similar markets have been used to study a variety of interesting topics. We believe that, by filling the gap between traditional experimental markets and the “real world,” markets such as the IEM have great research potential. The IEM includes more data than available from typical financial markets. This data includes transaction and order flow data associated with individual traders, complete queue information, portfolio positions of each trader and trader demographics. The IEM can also be used to survey traders at any time, recording survey responses and associating them with other data. Thus, they provide an excellent complement to other existing research techniques.

References

- Beckmann, K., and M. Werding (1996) “Passauer Wahlbörse’: Information processing in a political market experiment,” Kyklos, 49:171-204.
- Berg, J.E., R. Forsythe and T.A. Rietz (1996) “What makes markets predict well? Evidence from the Iowa Electronic Markets,” in W. Albers, W. Güth, P. Hammerstein, B. Moldovanu and E. van Damme (eds.), Understanding Strategic Interaction: Essays in Honor of Reinhard Selten, pp 444-463 (Springer, New York).
- Berg, J.E., R. Forsythe and T.A. Rietz (1997) “The Iowa Electronic Market,” in D. Paxson and D. Wood (eds) Blackwell Encyclopedic Dictionary of Finance (Blackwell, Oxford UK).
- Berg, J.E., F.D. Nelson and T.A. Rietz (2001) “Accuracy and forecast standard error of prediction markets,” mimeo, Department of Finance, University of Iowa.
- Berg, J.E. and T.A. Rietz, (2003) “Prediction markets as decision support systems,” Information Systems Frontiers, 5:79–93.
- Bondarenko, O., and P. Bossarts (2000) “Expectations and learning in Iowa,” Journal of Banking and Finance, 24:1535-1555.
- Forsythe, R., F.D. Nelson, G.R. Neumann and J. Wright (1991) “Forecasting elections: A market alternative to polls,” in TR Palfrey (ed.) Contemporary Laboratory Experiments in Political Economy, pp 69-111, (University of Michigan Press, Ann Arbor).
- Forsythe, R., F.D. Nelson, G.R. Neumann and J. Wright (1992) “Anatomy of an experimental political stock market,” American Economic Review, 82:1142-1161.
- Forsythe, R., T.A. Rietz and T.W. Ross (1999) “Wishes, expectations and actions: Price formation in election stock markets,” Journal of Economic Behavior and Organization, 39:83-110.
- Granberg, D. and E. Brent, (1983) “When prophecy bends: The preference-expectation link in U.S. Presidential elections,” Journal of Personality and Social Psychology, 45:477-491.
- Kahneman, D., and A. Tversky (1979) “Prospect theory: An analysis of decision under risk,” Econometrica, 47:263-291.
- Jacobsen, B., J. Potters, A. Schram, F. van Winden and J. Wit (2000) “(In)accuracy of a European political stock market: The influence of common value structures,” European Economic Review, 44:205-230.
- Oliven, K., and T.A. Rietz (2002) “Suckers are born but markets are made: Individual rationality, arbitrage and market efficiency on an electronic futures market,” mimeo, Department of Finance, University of Iowa.
- Ortner, G., (1997) “Forecasting markets: An industrial application: Part I,” mimeo, Technical University of Vienna.
- Ortner, G., (1998) “Forecasting markets: An industrial application: Part II,” mimeo, Technical University of Vienna.
- Plott, C.R., (2000) “Markets as information gathering tools,” Southern Economic Journal, 67, 2-15.
- Rietz, T.A. (1998) “Enforcing arbitrage restrictions in experimental asset markets,” mimeo, Department of Finance, University of Iowa.
- Slemrod, J., and T. Greimel (1999) “Did Steve Forbes scare the municipal bond market?” Journal of Public Economics, 74:81-96.

Figures

Trader: JDoe		Refresh		<i>Iowa Electronic Markets</i>		
US\$: 5.00000		11/01/00 11:59:59 PM		PRES00 VS		
Contract	BestBid	BestAsk	LastPrice	QuantityHeld	YourBids	YourAsks
<i>DemVS</i>	0.475	0.481	0.479	0	0	0
<i>ReformVS</i>	0.011	0.017	0.015	0	0	0
<i>RepVS</i>	0.510	0.518	0.514	0	0	0

-Market Orders- Qty Confirm? Market Order

-Limit Orders- Price Qty Expire Confirm? Limit Order

[PRES00_VS](#) [Go To Market](#) [My Account](#) [Logout](#)

Figure 1: IEM Market trading screen. The top table gives the current market, contracts, bids, asks and last trade prices along with the current trader's portfolio, number of outstanding bids and number of outstanding asks. The menu in the middle section allows the trader to place orders and make trades. The menu bar at the bottom allows the trader to undertake other activities.

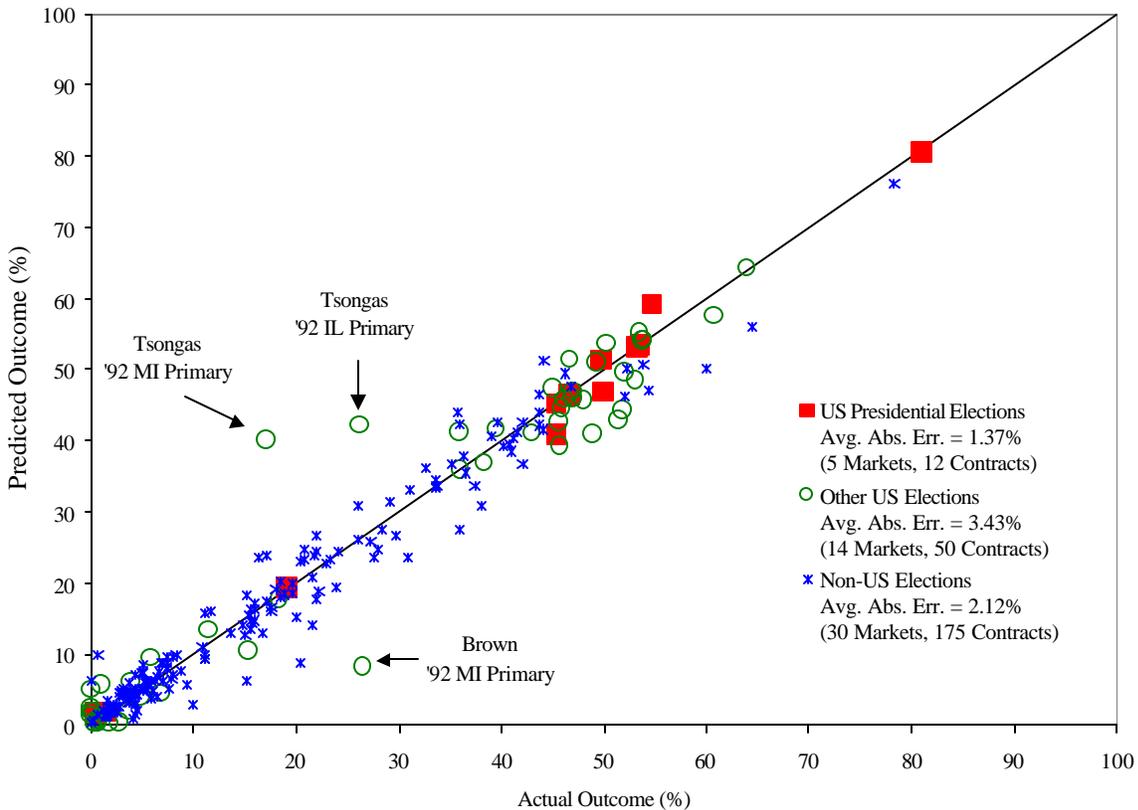


Figure 2: Political futures market predicted versus actual outcomes for vote-share and seat-share markets. Vote-share market predictions are for percentages of votes received by parties or candidates. Seat-share market predictions are for percentages of seats in congress or parliament held by parties. Predictions are based on normalized (to sum to 100%), last-trade prices as of midnight the night before each election. The 45-degree line represents perfect accuracy. Markets included in Figure 2:

Vote Share

Austria (Fed. Par. '95, Styrian Region '95; Vienna City '95, EU Par. '96)
 Canada (Par. '93, '96)
 Denmark (Par. '91)
 Finland (Pres. (x2 mkts.))
 France (Pres. '95)
 Korea (Pres. '92)
 Germany (Par. '90 (x3 mkts. in Bonn, Frankfurt and Leipzig), '91, '94, '98 Fed, '98 State)
 The Netherlands (Par. '91, '94 (2nd Chamber); EU Par. '94; Muni. Council '94)
 Norway (Par. '95)
 Sweden (EU Mbrshp. '94)
 Turkey (Par. '91)

Vote Share (Continued)

U.S. Pres. ('88, '92, '96 (x2 mkts.), '00)
 U.S. Pres. Primaries (IL '92, MI '92)
 U.S. Sen. (IL '90, IA '90, AZ '94, NJ '94, PA '94, TX '94, VA '94)
 U.S. Gov. (NY '94, TX '94)
 U.S. House (UT '94)

Seat Share & Other

Australia (Par. '93)
 Canada (Par. '93, '96)
 The Netherlands (Par. (2nd Chamber) Turnout '94, EU Par. Turnout '94)
 U.S. House ('94)
 U.S. Senate ('94)

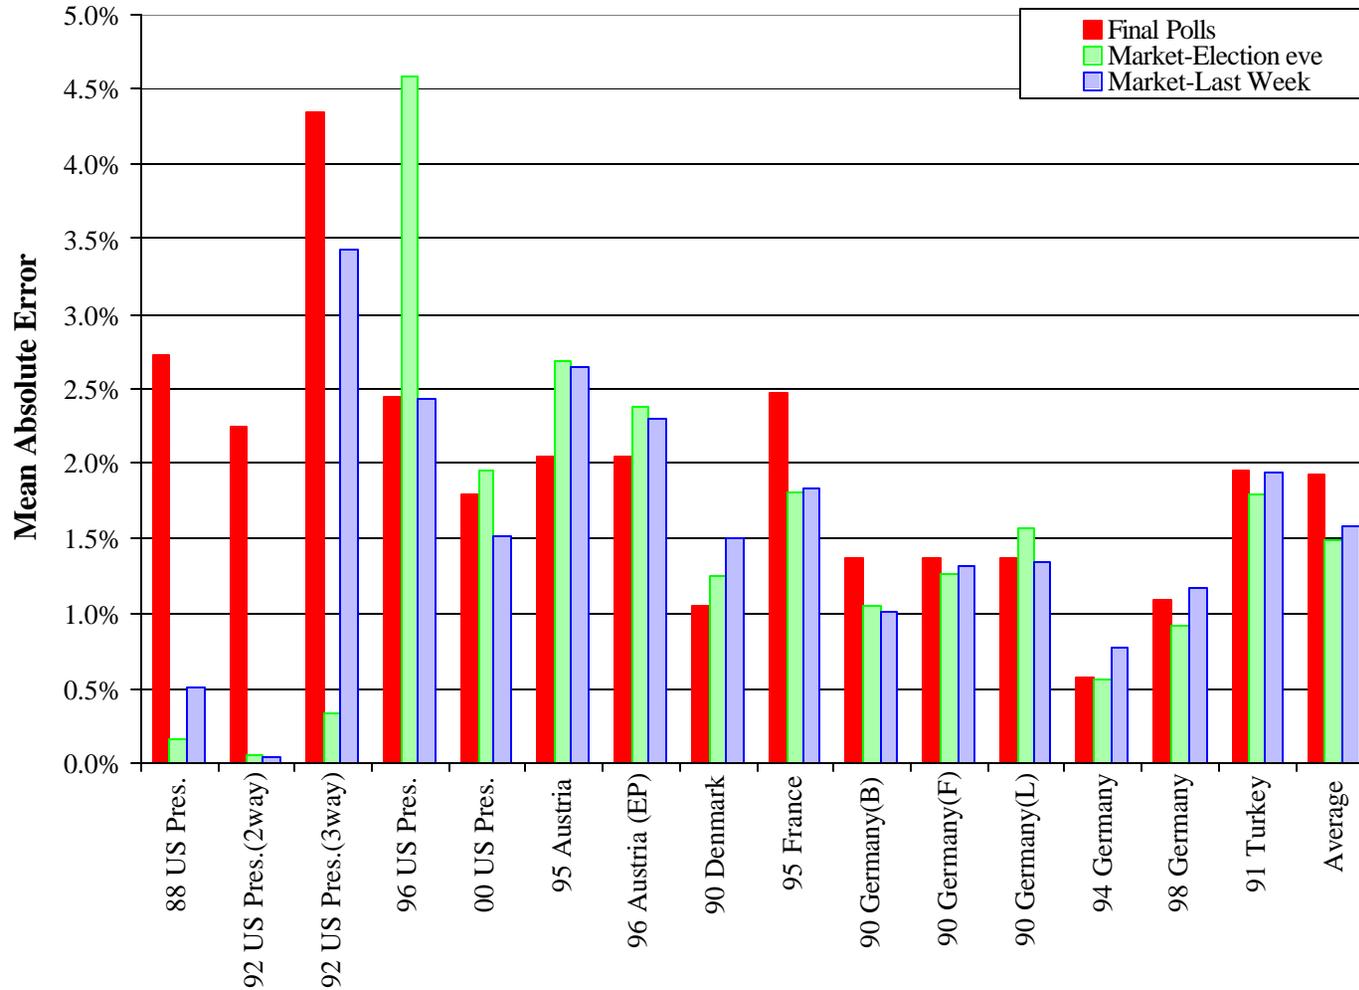


Figure 3: National political futures market average absolute prediction errors for candidate vote shares versus errors from polls in the week before the election. Final week market errors are calculated from the normalized (to sum to 100%) average trade prices from the last week before each election. Election eve market errors are calculated from normalized (to sum to 100%), last-trade prices as of midnight the night before each election. Errors for each poll are calculated from the normalized (to sum to 100% and allocate undecided voters) poll responses. Then, poll errors are averaged across known major polls from the last week before the election.

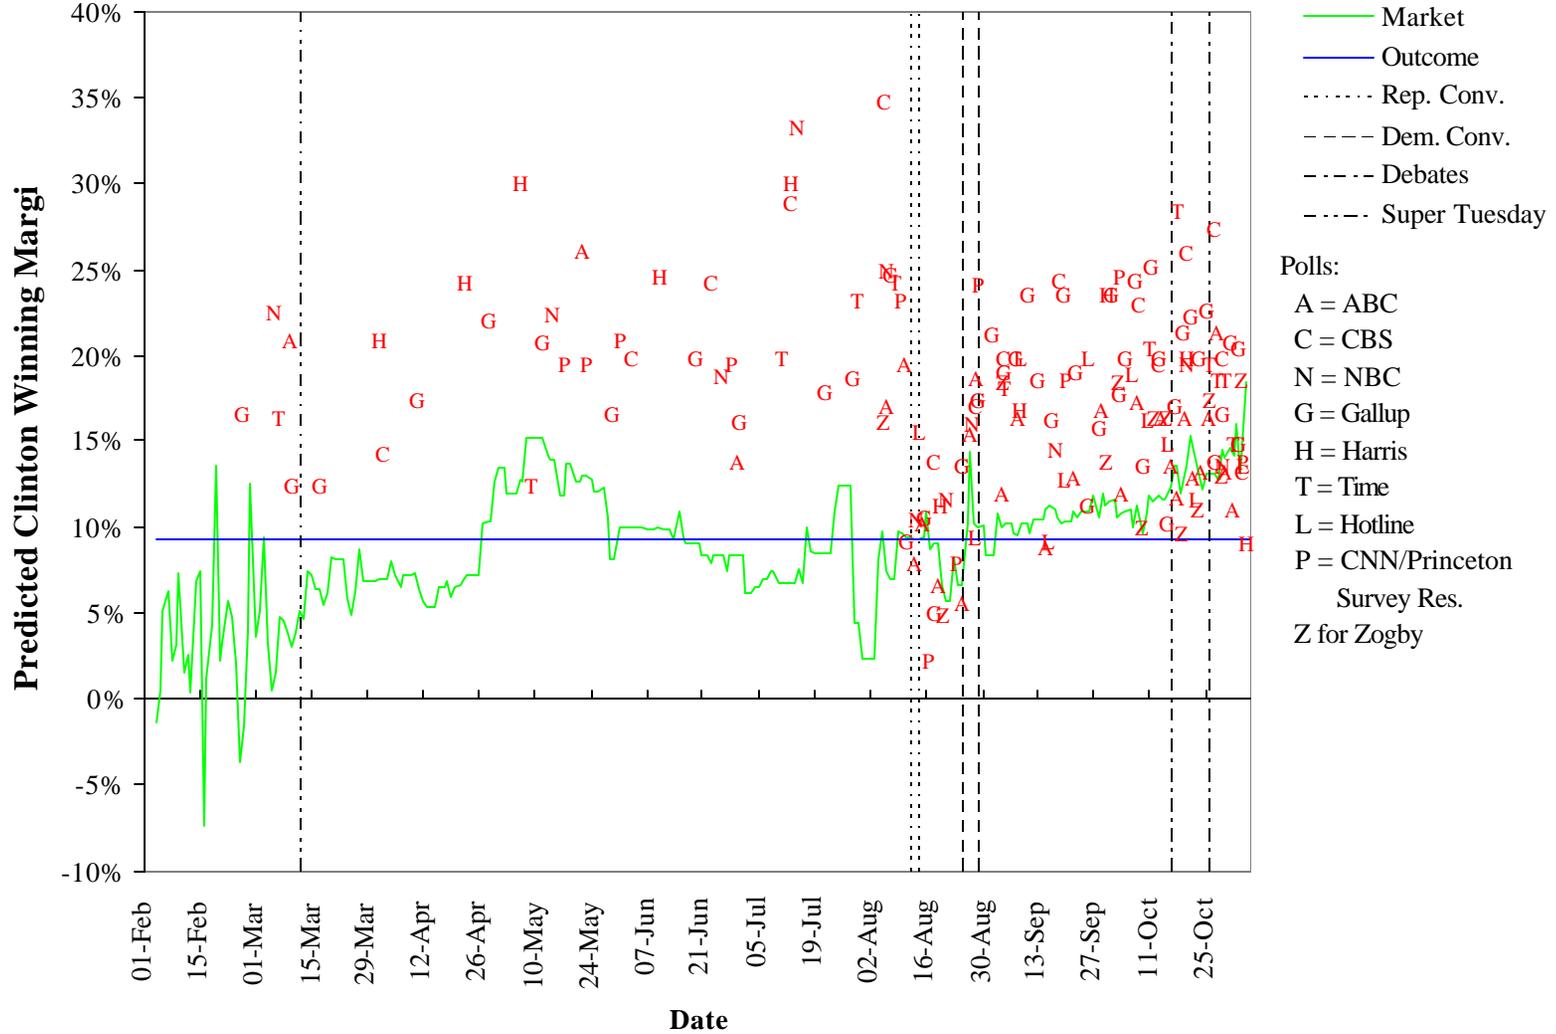


Figure 4: Performance of the 1996 Presidential Election Market and contemporaneous opinion polls over the full course of the campaign. The vertical axis is the normalized margin of victory for Clinton (Clinton outcome minus Dole outcome divided by the sum of the two.) Time is along the horizontal axis, the solid horizontal straight line indicates the outcome of the election on November 5, and the solid vertical lines represent important events – the Republican and Democratic conventions, two debates and Super Tuesday. The right-most limit of the graph is the election. The solid jagged line represents market prices, and the letters indicate the outcome of various polls.