Political Analysis (2004) 12:277–295
doi:10.1093/pan/mph019

Forecasting the Vote: A Theoretical Comparison of Election Markets and Public Opinion Polls

S. G. Kou
312 Mudd Building, Department of IEOR,
Columbia University, New York, NY 10027
e-mail: sk75@columbia.edu

Michael E. Sobel
421 Fayerweather Hall,
Columbia University, New York, NY 10027
e-mail: mes105@columbia.edu

The dominant methodology for short-term forecasting of electoral outcomes uses trial-heat polls, where respondents report their current electoral preferences (not their election-day predictions). Election markets, where self-selected participants trade shares of candidates at prices predictive of election-day results, provide an alternative method that often produces more accurate forecasts. Consequently, increasing attention is being paid to this methodology. However, it is poorly understood and lacks theoretical justification. Surprisingly, the rationale for forecasting using trial-heat polls has not been completely developed either. We develop the justification for using both election markets and public opinion polls to forecast electoral outcomes, giving conditions under which each method performs ideally. For the ideal case, we prove (under the reasonable assumption that market participants are aware of the poll results) that the mean square prediction error for the market forecast is smaller than that of any forecast based on one or more polls. The case in which the assumptions supporting each method fail is also considered. It is often reasonable to expect that the best case results hold approximately, in which case the market forecast should also beat any poll-based forecast. We also compare the bias and variance of market and poll-based forecasts; our results suggest the utility of using the series of market prices to study the course of campaigns.

1 Introduction

Political scientists have used preelection polls to study the (intended) voting behavior of different groups, the voting decision-making process, the impact of campaigns, and other phenomena. Though possibly of less interest to the majority of academic political scientists, polls (despite the fact that they technically elicit current electoral preferences) are also used to generate short-term election forecasts. Because these forecasts are readily

Authors' note: For helpful comments, we are grateful to David Austen-Smith, Mark Becker, Steven Durlauf, Andrew Gelman, Paul Glasserman, Brad Jones, Robert Shapiro, Christopher Wlezien, and several anonymous reviewers. This research is supported in part by NSF grants DMI-9908106 and DMS-0074637 and the John D. and Catherine T. MacArthur Foundation.

compared with election results, and predictions deemed too far off the mark reflect poorly on the responsible organization(s), survey research, and the media, accuracy has been a major concern in the public opinion community (for example, Mosteller et al. 1949; Crespi 1988; Lau 1994; Mitofsky 1998) for over 50 years.

During the past twenty years, political scientists have also generated a large literature on the determinants of electoral outcomes (Lewis-Beck and Stegmaier 2000). The models used there to predict previous electoral outcomes have also served as a basis for forecasting candidate shares of the vote in forthcoming elections (for example, Fair 1978, 1996; Rosenstone 1983; Erikson 1989; Beck 1992, 2000; Campbell 1992; Gelman and King 1993; Erikson and Wlezien 1994; Campbell and Garand 2000; Wlezien and Erikson, 2000). These longer-term forecasts, made several months to a year in advance of the election, often do not use trial-heat questions from preelection polls. Nevertheless, the better forecasting models generally yield more accurate predictions of the vote shares than public opinion polls conducted several months before the election (Gelman and King 1993).

While election forecasts derived from substantive considerations interest political scientists because the forecasts can be compared with electoral outcomes to test hypotheses, the more general subject of election forecasting is often regarded as tangential to the development of a science of politics, hence unworthy of serious professional attention. Such a view completely discounts the public need for accurate electoral forecasting.

Electoral outcomes in democratic countries have far-reaching domestic and (sometimes) international impact. Individuals, corporate actors, and governments who anticipate being affected by the outcome of a future election incorporate their expectations (forecasts) into current decisions and policies. Thus anticipated and actual electoral outcomes are associated with aggregate stock prices and foreign exchange rates (Blomberg and Hess 1997; Riley and Luksetich 1980; Brander 1991). Prices of selected stocks (or groups of stocks) may also be associated with anticipated outcomes (Herron et al. 1999; Knight 2003): for example, the price jump of pharmaceutical stocks one day before the 2000 U.S. presidential election, due to investors' beliefs that the Democratic party would not control both the executive and legislative branches of government (Kou and Sobel 2001). Major firms often deal with electoral uncertainty by making contributions to more than one candidate (party). A foreign government making an economic policy decision today may well take into account its expectations about the outcome of a forthcoming U.S. election. In all these instances, rational decision making hinges on an accurate assessment of uncertainty.

As opposed to poll-based forecasting, where the opinions of a probability sample of the electorate are aggregated, methods that combine the opinions of one or more experts are also sometimes used to produce election forecasts (see Rosenstone 1983 for a critical evaluation of many of these procedures). Election markets, in which participants buy and sell candidate shares, are a systematic variant of this approach. In "winner-take-all" markets, shares of the winning candidate pay out at positive values, while shares of the losing candidates pay out at 0. In vote share markets, a candidate's price is a forecast of his/her share of the vote. First implemented in large-scale fashion by researchers at the Univeristy of Iowa (Forsythe et al. 1992), election markets were used to predict the 1988 U.S. presidential election. On election eve 1988, the price of a Bush share was within one-tenth of a percentage point of the vote share Bush received. None of the major polls conducted over the last few days before the election did nearly as well. The market also predicted the marginalized lead (percentage difference between the top two candidates) much more accurately than any of the major polls. Subsequently, the Iowa Electronic Markets (IEM) have been used to study the 1992, 1996, and 2000 presidential elections, a number of local U.S. elections and primaries, and even foreign elections in Russia and France. Currently, anyone in the world who wishes to participate may do so. Similar markets have also been set up in other countries. Overall, markets, especially the larger ones, have tended to generate more accurate short-term forecasts than preelection polls (Berg et al. 1997; Forsythe et al. 1998, 1999; Bohm and Sonnegård 1999; see also Rhode and Strumpf 2003 for empirical evidences on historical presidential betting markets from 1868 to 1940). This suggests, if accurate short-term forecasting is the desideratum, that election markets merit serious attention.

Why, then, have political scientists devoted so little attention to these markets? First, unlike the polls, apparently the market data have not been widely perceived as useful for studying phenomena that are of theoretical interest. For exceptions, see Herron et al. (1999) and Shaw and Roberts (2000). Herron et al. (1999) use data from the 1992 U.S. presidential winner-take-all market to study the effect of election forecasts on equity prices in 74 economic sectors. Shaw and Roberts (2000) argue that the market prices better encode the information in campaign events than the polls; they therefore suggest using the market series of prices for studying the course of election campaigns. Second, the empirical results above are puzzling and lack theoretical justification. Whereas mathematical laws of large numbers and central limit theorems would appear to provide theoretical justification for polling, suggesting that in a large probabilistic sample of the electorate, taken close enough to election day for preferences to have stabilized, the sample percentage should be close to the true population percentage, a similarly compelling justification has not been provided for the use of election markets.

To be sure, various arguments have been offered ex post facto to explain the relative performance of polls and markets, with researchers pointing to various problems associated with polling and/or arguing that markets are "efficient" and hence forecast electoral outcomes well (for example, Forsythe et al. 1992; Berg et al. 1997). But such types of arguments provide no theoretical justification (and no real comfort) for using market forecasts. In the first case, the market forecast is merely empirically less "bad" than the poll forecast, hardly a theoretical argument for the use of markets. The second argument is vacuous on two counts. First, to claim that the market is efficient is merely to say that the market price is an efficient price. But this does not specify the price itself, which depends on the information set used by traders. (Further, in real markets, efficient prices are risk-adjusted and such prices would not predict electoral outcomes well. So even the intuition that an efficient price should predict well is in general incorrect.) Second, no formal criteria have been put forth for comparing the performance of markets and polls. In fact, as we later demonstrate, even if the market price is efficient, it does not follow that the market forecast should theoretically dominate predictions made using a poll (or set of polls).

In this article we make several contributions. First, we provide a theoretical justification for using election markets to forecast vote shares and winners in upcoming elections. (More generally, our results provide theoretical justification for using prices to aggregagte information in experimental markets.) Second, we develop the rationale for using poll-based forecasts (surprisingly, this has not been developed in previous work). Third, we propose criteria for comparing different forecasts. Fourth, we provide conditions under which market forecasts should be preferred to poll-based forecasts (whether the latter are based on one poll or many polls). Fifth, our comparison of the biases and variances of these forecasting methods suggests the potential utility of using markets to study campaigns. Sixth, we consider the case in which one or more assumptions supporting the use of markets or polls fail and we assess the likely impact on the results. Whereas previous discussions of such matters have been ad hoc, because we have theoretical results as a baseline we can discuss this issue systematically and rigorously.

We proceed as follows. Election markets are introduced in the next section. In Section 3, equilibrium candidate prices (market forecasts) are derived and conditions under which the market trades at these prices are given. Next, systematic and random sources of error in the polls are considered. Analytical results on forecast errors from polls and markets are then provided. The two methods are compared both when the theoretical justifications underlying each approach hold and when these theoretical justifications do not hold.

2 Polls and Election Markets

In this section, election markets are introduced and the market series is informally compared with the Gallup/CNN/USA Today poll series for the 2000 U.S. presidential election. The intent here is only to introduce the markets to readers and to highlight a few points that emerge from looking at the two series, not to conduct a full-scale empirical comparison.

The IEM U.S. presidential vote share market for the 2000, election opened on January 5, 2000 and closed November 10, 2000, three days after the November 7 election. Three parties (Democrat, Republican, Reform) were featured.[1] The market prices forecast the share of the vote that each of the three parties would receive on election day. Figure 1 is a graph of the daily closing prices in this market. Figure 2 is a histogram of empirical vote shares (of the three-party vote) for the Democratic, Republican, and Reform candidates using the series of Gallup/CNN/USA Today polls conducted between April 28–30, 2000, and August 27, 2000.[2] Unlike the market forecast, the polls solicit the current electoral preferences of respondents. Figure 3 reports the empirical vote shares from September 4 to November 6, the day before the election.[3] The estimates are based on the responses of "likely voters," that is, the subset of respondents who report being registered to vote and who Gallup deems likely to vote, based on their responses to questions regarding previous voting behavior and current voting intentions.[4]

[1]The market organizers did not anticipate the Green party candidacy nor include a residual category for other candidates, but liquidation values (the price paid to shareholders at the close of the market) were defined as percentages of the three-party vote share (for example, percentage voting Democratic among those voting for the Democratic, Republican, or Reform party).

[2]Interviewers provided respondents with the names (as soon as these were available) and affiliations of the candidates (George W. Bush and Dick Cheney—Republican; Al Gore and Joe Lieberman—Democrat; Pat Buchanan and Ezola Foster—Reform; Ralph Nader and Winona LaDuke—Green) and asked the respondents who they would vote for "if the election were being held today." Undecided respondents who reported leaning to one of the above categories were assigned to that category; the remainder of the undecided were coded as undecided. This yields a percentage distribution over five categories. To render the polling data comparable to the market data, for each poll we computed the proportion p of undecided or Green responses, reweighting the other three proportions by $(1 - p)^{-1}$. Our estimates implicitly allocate the undecided voters who do not eventually opt for the Green party among the three parties in the proportions above. If this is not the correct allocation, bias is introduced. However, the amount of bias cannot be substantial, because the undecided never comprise more than 8% (poll reported for October 9–11) and more typically comprise about 5%, declining to 3% in November.

[3]Interviewing was conducted daily from September 4 through November 6, and until November 4 results were reported on the basis of three-day rolling averages (September 4–6, September 5–7, etc.). The data from November 4 and 5 were also pooled, as were the data for November 5 and 6.

[4]The practice of using "likely voters" in estimating vote shares is common throughout the industry, with different organizations employing different (usually proprietary) criteria for deciding who to include (exclude) in the calculations. For further description and criticism of Gallup's methodology, see Erikson (1993). The number of likely voters on which the estimates in Fig. 2 are based range from 499 (April 28–30) to 697 (August 18–19). For Fig. 3, the numbers range from 621 (September 26–28) to 851 (October 24–26), increasing after October 24–26, with more than 2000 likely voters in each of the polls conducted the week prior to the election.

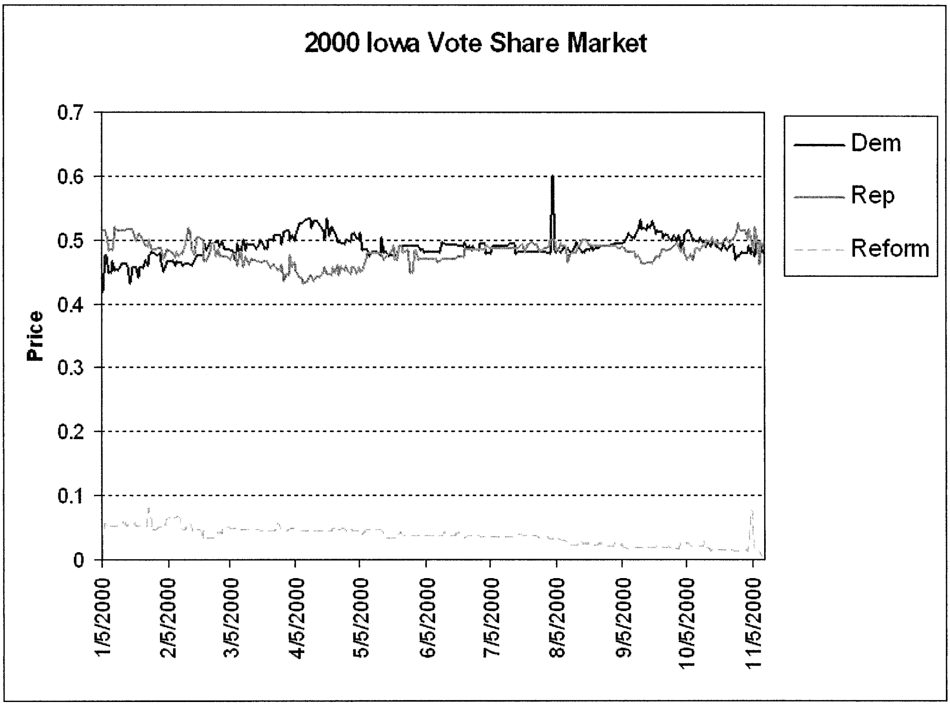

Fig. 1 Iowa Electronic Market between January 5, 2000, and November 7, 2000.

Prior to the Republican convention (July 31–August 3), the polls showed the Republicans leading the Democrats, typically by five or more percentage points, with a difference of 12 points in the poll of July 25–26. After the convention, the lead increased to 18 percentage points, remaining at that level through the poll of August 11–12. The Democratic convention was held August 14–17, and in the next poll (August 18–19), the Democrats pulled ahead by one percentage point. This lead appeared to widen until the third week of September, narrowing thereafter and reappearing briefly around the beginning of October, with the Republicans ahead thereafter. Note that the final polls (with larger samples) showed the gap narrowing.

By way of contrast, with the exception of the one-day spike in the Democratic price after the Republican convention, which is due to the very thin volume on that date and irregular trading, the market projection suggests a much closer race throughout. The market data do not exhibit the volatility (or the "convention effects") typical of data from the polls, suggesting that market participants are able to separate transitory and permanent shifts in the preferences of the electorate, effectively discounting the former when making the market forecast. This fits well with the theoretical model in the next section.

Turning to the last few days of the election cycle, the November 5–6 poll vote shares were .48 for the Democratic ticket, .51 for the Republican ticket, and .01 for the Reform ticket. For November 5, the market forecasts (the closing prices) were .498 (Democrats), .485 (Republicans), and .017 (Reform), and for November 6 were .475 (Democrats), .520 (Republicans), and .017 (Reform). The liquidation values (vote shares reported in the *New York Times* on November 10) were .499 (Democrats), .497 (Republicans), and .004 (Reform). Note that both the polls and the market overpredicted the Reform share and underpredicted the Democratic vote share. The market projections for the Democratic and

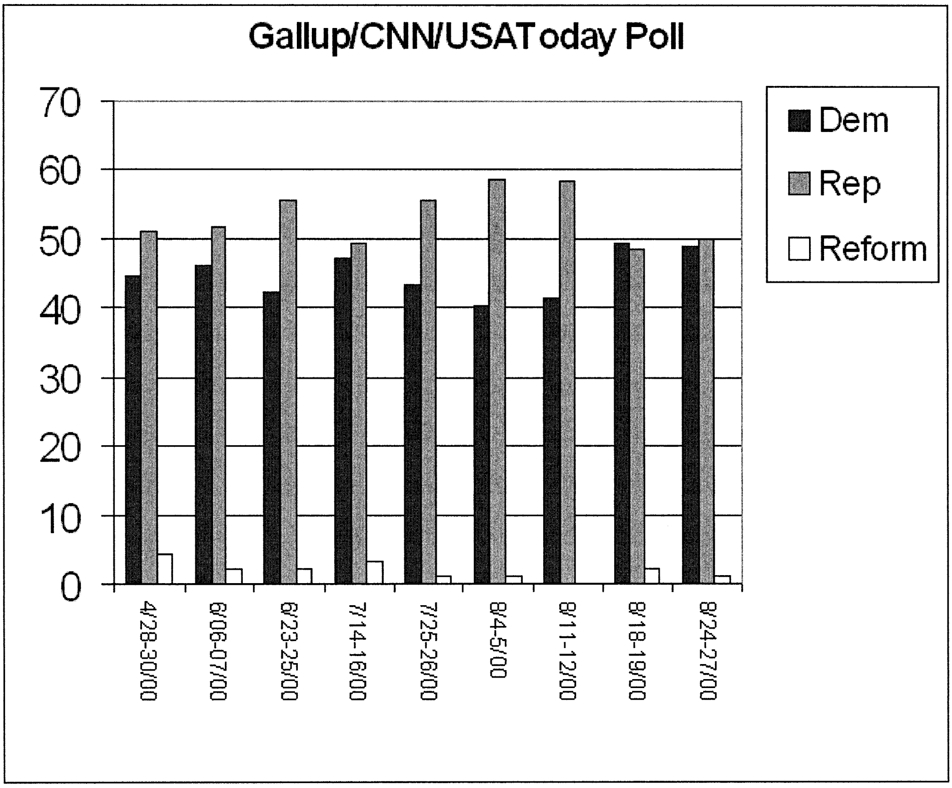

Fig. 2 Gallup/CNN/USA Today polls between April 28–30, 2000, and August 27, 2000.

Republican vote shares are closer to the mark than the poll results on November 5, but not on November 6. Although both the market and poll predictions (at least in this poll) were close to the mark in the 2000 election, note also that the election was so close that on November 6, the polls predicted that Bush would win the popular vote (as did the market on November 6, but not on November 5).

3 Pricing the Candidates: The Market Forecast

3.1 *Motivation*

The ex post facto efficiency "explanation" for why markets forecast well argues that the market price is a good predictor because the market is efficient (Forsythe et al. 1992; Berg et al. 1997). But because the argument fails to define an efficient market (or an efficient price), efficiency is merely a synonym (without explanatory force) for predictability. If the efficiency argument is to have explanatory force, an efficient market (efficient prices) must be defined independently of its predictive ability.

One strategy is to define efficient prices and simply hypothesize that election markets are efficient. There are two problems here. First, this fails to provide any insight into why election markets work. It is certainly not intuitively obvious why the market should produce efficient prices that turn out to be good forecasts. In fact, standard financial theory would suggest that the equilibrium prices are given under "risk-adjusted" probabilities

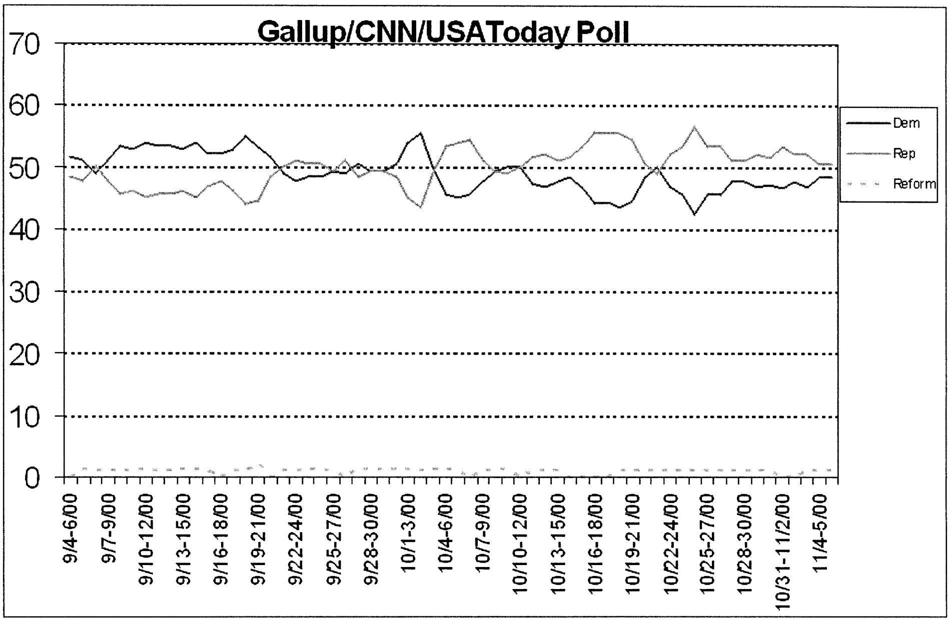

Fig. 3 Gallup/CNN/USA Today polls between September 4–6, 2000, and November 5–6, 2000.

(Hull 2000), which are very different from the real-world probabilities that are needed for forecasting. Second, without some understanding of conditions under which election markets (and polls) forecast well, little can be said about their relative performance in the real world, where one or more of these conditions might fail.

Thus, to put the efficiency argument on firm ground, equilibrium (efficient) candidate prices in the IEM markets must first be defined (independently of the market price) and conditions under which such prices obtain must be given. In this section, equilibrium prices are derived from first principles, using several standard types of assumptions about the behavior of market participants and several distinctive features of the IEM (features not shared by real markets). Of critical importance in the latter regard is the investment cap of $500 in the IEM. Due (in part) to this small cap, the IEM is not federally regulated. But more important, it is thereby reasonable to assume that a participant's investment in the IEM vote share market is not part of his/her overall investment strategy.

Letting T index the time of the election, we note that poll results at time $t < T$ elicit the current (time t versus time T) preferences of respondents. By way of contrast, the key result we derive in this section is that the equilibrium price of a candidate at time $t < T$ in the vote share market is his/her expected vote share at time T (conditional on the information set used by the market participants at time t). We emphasize this result because our focus is on comparing vote share markets with polls. However (as will be evident subsequently), our results also pertain to the winner-take-all markets. That is, for this case the equilibrium market price of a candidate at time $t < T$ is that candidate's probability of winning the election at time T.

Our results are neither obvious nor or a priori intuitive, for in the real-world futures markets studied in finance and economics, which are used by investors for hedging and speculating, the equilibrium prices are risk adjusted. Such prices, which would differ from those obtained herein, would not be useful for forecasting electoral outcomes (Kou and Sobel 2001).

More generally (as documented in Pennock et al. 2001), experimental markets have proven useful for predicting such disparate phenomena as Hollywood box office receipts and Oscar winners (Hollywood Stock Exchange), the timing of scientific and technological breakthroughs (the Foresight Exchange), and the results of Formula One racing. Similarly, markets have been used to predict outcomes of World Cup soccer matches and the timing of future events within organizations. Experimental markets have also proved useful in aggregating information in a number of other contexts, for example, in market research (Chan et al. 2001; Spann and Skiera 2003). In that vein, our results also apply to these types of markets and thus provide a very general rationale for the use of experimental markets.

3.2 *A Theoretical Justification*

We begin by describing the organization of the market. Anyone may trade by setting up an account with the market organizers and buying \$1 unit portfolios. Each unit portfolio consists of one share of each featured candidate (or party). Shares are traded on-line in a "double auction market" (both buyers and sellers present bids) and/or additional purchases are made using funds in the account, as in regular electronic stock markets such as NASDAQ. At any time, a participant may buy or sell unit portfolios for \$1, effectively making the market rate of interest $r = 0$. Buying on margin and short selling are not permitted, but an artificial short position can be created by purchasing unit portfolios and selling the shares of the candidate one wants to short.[5] The market is closed after the election. In the vote share market, the proportion of the vote received by the candidate (among those featured) is the value at which shares of the candidate are liquidated (paid out). In the winner-take-all market, shares of the winning candidate pay out at \$1 and shares of all the other candidates pay out at \$0. For further details on the structure of these markets, see Forsythe et al. (1999).

Suppose there are $c = 1, \ldots, C$ candidates featured in the election market and the market opens at time $t = 0$; the election is held at time T. For example, in the IEM vote share market previously discussed for the 2000 U.S. presidential election, $C = 3$ (Democratic candidate, Republican candidate, Reform candidate), and $t = 0$ corresponds to the opening of the market on January 5, 2000. The proportion of the vote (among candidates $c = 1, \ldots, C$) that c would receive if the election were held at time t is denoted by $\theta_c(t)$; this "latent" variable is observed only at time T.

To derive equilibrium candidate prices, several assumptions about the structure of the market and the behavior of participants are needed. A fuller discussion of the reasonableness of these assumptions is deferred to the section on comparing the markets and polls.

To begin, we assume:

M1: Investments in the election market are not part of a participant's overall investment strategy.

[5] Buying on margin is essentially borrowing money from a broker to purchase shares (or futures contracts, and the like). To cover the loan, the investor sets up a margin account with the broker, depositing some percentage of the price of the shares. Each day the investor's gains and losses are tabulated and added or subtracted from the balance in the margin account. If the balance dips below a certain point, the broker issues a margin call to the investor. The investor must then either replenish the margin account or the broker closes out the account, selling the shares at a loss. Short selling is the selling of shares (or other assets) borrowed from the broker. As above, the investor sets up a margin account, and if the asset price rises enough that the margin account dips below a certain point, a margin call is issued. The investor must then replenish the account or the broker closes out the account.

Assumption M1 is key and highlights two features of the election market (and other experimental markets) not shared by real financial markets. It is reasonable on several grounds. First, investments are limited to \$500 per participant. Second, by investing in the general asset market, an investor can expect to earn at least the risk-free interest rate $r > 0$.[6] By way of contrast, the interest rate in the IEM is 0.

At time 0, consider a participant with funds $x \leq \$500$ who trades shares at $P_c(0)$ per share, purchasing net amounts $\delta_c \geq 0$ of each candidate (subject to the constraint $\sum_{c=1}^{C} \delta_c \leq x$), yielding wealth

$$W(0) = \left(x - \sum_{c=1}^{C} \delta_c \right) + \sum_{c=1}^{C} \frac{\delta_c}{P_c(0)} \theta_c(T) \tag{1}$$

at time T, with utility $U(W(0))$. Note that if (M1) were false, $W(0)$ should be redefined to also include the participant's investments in other assets. We also assume:

M2: The utility function U is continuously differentiable and $U' \neq 0$.

The participant, who does not know $\theta_c(T) \equiv P_c(T)$ prior to T, uses information $\mathcal{F}_0$ to form the subjective expected utility $E_s(U(W(0)) \mid \mathcal{F}_0)$. Intuitively, when the candidate shares are priced so that the participant is indifferent between trading and not trading in the market, an equilibrium has been reached. More precisely, let $\delta = (\delta_1, \ldots, \delta_C)$ and define the equilibrium price vector as the set of prices $(\tilde{P}_1(0) > 0, \tilde{P}_2(0) > 0, \ldots, \tilde{P}_C(0) > 0)$ such that $\frac{\partial^+(E_s(U(W(0)) \mid \mathcal{F}_0))}{\partial \delta} \mid_{\delta=0} = 0$ holds for all x and utility functions U, where $\frac{\partial^+}{\partial \delta}$ denotes the right-hand derivative. Indeed, if the cth component is positive, then the participant has an incentive to buy shares of candidate c, and if the cth component is negative, the participant has an incentive to buy unit portfolios and then sell shares of candidate c. Either way, an equilibrium has not yet been reached. This definition is adopted from the literature on options pricing and rational expectations (Lucas 1978; Karatzas and Kou 1996; Davis 1997).

Using M2, it follows that

$$\begin{aligned} 0 &= \frac{\partial^+(E_s(U(W(0)) \mid \mathcal{F}_0)}{\partial \delta_c} \\ &= E_s\left(U'\left(x - \sum_{c=1}^{C} \delta_c + \sum_{c=1}^{C} \frac{\delta_c}{\tilde{P}_c(0)} \theta_c(T) \right) \left(-1 + \frac{\theta_c(T)}{\tilde{P}_c(0)} \right) \mid \mathcal{F}_0 \right), \end{aligned} \tag{2}$$

and setting δ to 0 gives $\tilde{P}_c(0) = E_s(\theta_c(T) \mid \mathcal{F}_0)$; that is, the time 0 equilibrium price for any candidate c is the subjective conditional expectation of c's vote share on election day. Note that the equilibrium price does not depend on the specific form of the utility function. In particular, we do not even have to assume that market participants are risk averse.

The definition of the time 0 equilibrium price and the conclusions above are readily extended to find the time t equilibrium price $\tilde{P}_c(t) = E_s(\theta_c(T) \mid \mathcal{F}_t)$ for any time $t \in [0, T)$ and information set $\mathcal{F}_t$. Note that the information set $\mathcal{F}_t$ will be increasing as t increases, $\theta_c(t) \notin \mathcal{F}_t$ for $t < T$ [as $\theta_c(t)$ is not observed for $t < T$], $\theta_c(T) \in \mathcal{F}_T$, and $\mathcal{F}_{T_-} \neq \mathcal{F}_T$.

[6]The risk-free rate is the theoretical interest rate that is paid out on an investment with no risk. In practice, the rate (in the United States) is typically taken to be the rate on a short-term treasury note, the assumption being that the government will honor its domestic financial obligations with certainty because it can issue the currency.

M3: Participants have access to the same information set $\mathcal{F}_t$ at time t and form identical expectations.

With M3, $E_s(\theta_c(T) \mid \mathcal{F}_t)$ is now the equilibrium price for all participants. Relaxation of M3 will be discussed later.

To insure that shares trade at equilibrium candidate prices, we make two assumptions:

M4: The market is in equilibrium.

M5: (A) No one participant can influence the market. (B) There are no transaction costs involved in trading. (C) The bid-ask spread is 0.[7]

M5 (or some stronger assumption) is usually made in analyzing financial markets, not because it is literally true, but because as long as it holds approximately, as in large markets, analysis is simplified with little loss. Finally, we assume:

M6: Expectations are rational.

This assumption, standard in economics and finance (Lucas 1978; Stokey and Lucas 1993), allows us to equate the objective (real-world) expectations $E(\theta_c(T) \mid \mathcal{F}_t)$ with $E_s(\theta_c(T) \mid \mathcal{F}_t)$.

Under M1–M6, for any candidate c, the market price is $\tilde{P}_c(t) = E(\theta_c(T) \mid \mathcal{F}_t)$. The market uses the current available information $\mathcal{F}_t$ to forecast the election day outcome $\theta_c(T)$ (not the current unobserved state $\theta_c(t)$ that is reflected by the polls). This helps to account for the different reaction of the markets and the polls to the conventions, as revealed in Figs. 1–3.

Because $\sum_{c=1}^{C} \theta_c(T) = 1$ (the vote shares must add to 1) and the expectation operator E is linear, $\sum_{c=1}^{C} \tilde{P}_c(t) = 1$, that is, the equilibrim prices add to 1, meaning participants cannot profit through arbitrage. This should be true even if the market is not in equilibrium, for if the prices add to more (less) than 1, a participant can buy (sell) unit portfolios and sell (buy) the individual shares, thereby locking in a profit at no risk. Typically, such arbitrage opportunities, if they appear, also disappear rapidly.

It is important to note that the equilibrium prices we have obtained cannot be derived from the standard approaches (the capital asset pricing model, arbitrage pricing theory) to the valuation of financial derivatives. First, in these approaches there is an asset (for example, a stock) with an observed current value $S(t)$ underlying the financial derivative, and $S(t)$ is assumed to follow the log-normal distribution. However, the underlying asset in the vote share market is the percentage $\theta_c(t)$, which cannot possibly follow the log-normal distribution and is also unobserved. Second, empirical work suggests a correlation between election results and the currency and stock markets (see, for example, Riley and Luksetich 1980; Blomberg and Hess 1997). In this case (leaving aside the first problem just mentioned), routinely applying these approaches would yield equilibrium prices $E^*(\theta_c(T) \mid \mathcal{F}_t) \neq E(\theta_c(T) \mid \mathcal{F}_t)$, where E^* is a risk-adjusted expectation, possibly dependent on the strength of this correlation. However, our derivation is not affected by this correlation because assumption (M1) decouples investments in the election market from investments in the general asset market. (Recall that this assumption was motivated by two unusual features of the IEM not shared by real markets, namely, the 0 interest rate and the investment cap of \$500.)

It follows from the above that to apply one of the standard pricing approaches, it is necessary to make the additional assumption M1 (or a very similar assumption) and to also make an extra assumption that some function of the vote shares followed

[7]The ask price is the (higher) price the seller is asking while the bid price is the (lower) price that the buyer is offering. The bid-ask spread is the difference between the two prices.

the log-normal distribution. Our approach is both more parsimonious and more straightforward.

Our analysis also reveals that were M1 to fail, as might happen were the $500 cap removed, real investors might use the market for hedging, in which case the market prices would no longer necessarily be useful for forecasting. For further consideration of alternatively structured electoral markets useful for hedging the financial risk associated with electoral outcomes, see Kou and Sobel (2001).

The derivation above pertains to the vote share market because the liquidation price $P_c(T) \equiv \theta_c(T)$ of candidate c is defined to be candidate c's vote share on election day. Had $P_c(T)$ been defined as 1 if candidate c wins the election, 0 otherwise, as in winner-take-all markets, the same argument would lead to equilibrium prices $\tilde{P}_c(t)$ equal to the probability at time t that candidate c will win the election at time T. Thus our results also provide theoretical justification for interpreting the prices in winner-take-all markets as predicted probabilities of winning and for using these markets to forecast election winners and losers.

Finally, our results explain why prices are useful for predicting and aggregating information in similarly structured experimental markets, for example, the Hollywood Stock Exchange. As above, assumption M1 is critical; without this assumption, there is no reason to believe that the market prices would have predictive value.

4 Polling the Voters: The Sampling Approach

Preelection polls elicit current electoral preferences at times $t < T$. The vote share $V_c(t)$ obtained by aggregating the responses to a particular poll at time t (possibly combined with other information from that poll) is typically regarded not as a forecast of $\theta_c(T)$ but rather as an estimate (with error) of the latent vote share $\theta_c(t)$ that would be observed if the election were held at time t:

$$V_c(t) = \theta_c(t) + B(\theta_c(t)) + \varepsilon_c(t), \tag{3}$$

where $B(\theta_c(t)) = E(V_c(t) - \theta_c(t) \mid \theta_c(t))$ is the systematic measurement error (bias), and $\varepsilon_c(t)$, with $E(\varepsilon_c(t) \mid \theta_c(t)) = 0$, is the random component of error. In a large sample, $\varepsilon_c(t)$ should converge in probability to 0; if in addition $B(\theta_c(t)) = 0$ (or converges in probability to 0 in large samples), $V_c(t)$ will be a good estimate of $\theta_c(t)$.

When $V_c(t)$ is also used to forecast $\theta_c(T)$, as is typical when t is close to T, an additional source of error, namely, the *projection error* $\theta_c(T) - \theta_c(t)$, is introduced. When t is not close to T (for example, several months before the election), many would argue that the probability of a large projection error is too high for it to make sense to use $V_c(t)$ as a forecast. However, if t is close to T, if voters' preferences have hardened, the projection error should be small, hence $V_c(t)$ should also constitute a useful forecast. Thus polls that elicit current personal preferences (as versus, for example, asking respondents who they think will win) may be useful for short-term forecasting but not generally for long-term forecasting.

This section focuses on the random and systematic measurement error in using $V_c(t)$ to estimate $\theta_c(t)$ for $t < T$ and gives conditions under which the systematic measurement error vanishes. This is the best-case scenario for a poll. In stating these conditions, we are not arguing that they are actually met in practice; this issue is addressed in Section 5.2. We do not consider the projection error in this section. That subject is addressed extensively in the next section, where the equilibrium prices $\tilde{P}_c(t)$ are compared with vote shares $V_c(t)$ (both when there is systematic bias and when there is not).

To begin, hundreds of organizations conduct preelection polls. Our concern is with only high-quality polls using probabilistic sampling, such as those conducted by the major news media. Thus we assume:

S1: A probabilistic sample of the eligible voting population has been taken.

Some of the sampled are never contacted and others refuse to participate. If response rates and voting preferences vary systematically with demographic characteristics of the population, the failure to take this into consideration can lead to poor estimates of $\theta_c(t)$. Polling organizations therefore use poststratification to adjust the distribution of demographic characteristics among respondents to the distribution in the population (Voss et al. 1995). The adjustments assume that nonrespondents within demographic categories have the same preferences as the respondents in these categories, that is, the missing data are missing at random (Little and Rubin 1987). If this is not correct, systematic bias may be induced. Thus we explicitly assume away this source of bias:

S2: Estimates adjusted for nonresponse are free of systematic measurement error.

Even if S1 and S2 hold, only some of those eligible vote. Since it is not reasonable to assume that the preferences of voters and nonvoters are identical, most organizations discard, do not solicit, or adjust in some other fashion for the preferences of respondents deemed unlikely to vote (Crespi 1988; Voss et al. 1995). To remove this source of bias, we assume:

S3: Estimates adjusted for likely voting status are free of systematic measurement error.

In addition, some respondents report being undecided. If interest focuses on using the polls for forecasting (or even for predicting the results of an election held at time $t < T$), this can be problematic. Especially as the election nears, a number of the major polling organizations apportion the percentage undecided (or the percentage of the undecided deemed likely to vote, for example) among the candidates (Crespi 1988; Newport 1997; Panagakis 1997; Traugott and Lavrakas 2000). Different assumptions about the behavior of the undecided lead to different allocation rules. A common procedure allocates the percent undecided to the two major candidates, in proportion to the estimated share of the vote currently held by each. We assume:

S4: The procedure used to apportion the votes of the undecided voters is free of systematic measurement error.

Although other potential sources of systematic measurement error have been identified, those above are the most substantial. For example, although respondents have no incentive to be truthful (Forsythe et al. 1992), they typically have no reason to be dishonest, either. Variations in methods and questions used to ascertain preferences may also make a difference, though the evidence in Gelman and King (1993) suggests that the impact of this is minor. Further, in the high-quality polls, these sources are unlikely to produce systematic measurement error. Nevertheless, to be explicit, we assume:

S5: There are no other sources of systematic measurement error.

Under S1–S5, $B(\theta_c(t)) = 0$; consequently $E(V_c(t) \mid \theta_c(t)) = \theta_c(t)$.

Of course, whether or not S1–S5 hold, the discrepancy between $\theta_c(T)$ and $\theta_c(t)$ (the projection error), discussed in the next section, must be also considered when $V_c(t)$ is used for forecasting:

$$\theta_c(T) - V_c(t) = (\theta_c(T) - \theta_c(t)) - B(\theta_c(t)) - \varepsilon_c(t). \tag{4}$$

5 Comparing Markets and Polls

5.1 *Theory*

We propose an explicit criterion (the mean square prediction error, hereafter MSPE) for comparing alternative forecasts $\hat{\theta}_c(T)$ of $\theta_c(T)$ with respect to their predictive ability:

$$MSPE(\hat{\theta}_c(T)) \equiv E(\hat{\theta}_c(T) - \theta_c(T))^2. \quad (5)$$

(The MSPE should not be confused with the more familiar mean square error [MSE].) In this section we focus on the use of $V_c(t)$ as a forecast of $\theta_c(T)$, versus an estimate of $\theta_c(t)$ (as in the previous section when we studied polls). We then compare MSPE($\tilde{P}_c(t)$) with MSPE($V_c(t)$). In general, there is no relationship between these two quantities, implying that neither forecast dominates the other. However, in the important special case in which the information used by market participants includes the results from the poll compared with the market, $MSPE(\tilde{P}_c(t)) \leq MSPE(V_c(t))$, that is, the equilibrium price dominates the vote share $V_c(t)$, regardless of systematic bias in the vote share. We also prove a much stronger result. Suppose the information used by market participants includes the results from polls $g = 1, \ldots, G$. The equilibrium price then dominates any (measurable) function of the poll results (for example, the average of all the poll results). In other words, the market does better than any scientist can do using the results of these polls.

Shaw and Roberts (2000) argue that the market prices better encode the information in campaign events than the polls, suggesting the utility of using the series of prices for studying the course of election campaigns. We prove two theorems that jointly offer support for this suggestion. First, we show that the prediction bias of the equilibrium price is 0, whereas this is not generally true of the polls (whether or not the systematic bias vanishes). Next, we give reasonable conditions under which the variance of the equilibrium price is smaller than the variance of the vote share $V_c(t)$; this result may account for the fact that the market appears less volatile than the polls. Finally, we consider MSPE($V_c(t)$) at greater length, focusing on the behavior of this forecast as election day nears. We further give a set of conditions under which the public opinion polls will yield a good forecast as election day nears. Although conventional wisdom suggests this to be the usual case (were this not the case there would be no point in using polls to forecast election results), this issue has not been treated carefully in the literature.

Throughout (until the subsection on practice) we assume that M1–M6 hold, implying that the market equilibrium price $\tilde{P}_c(t) = E(\theta_c(T) \mid \mathcal{F}_t)$. That is, $\tilde{P}_c(t)$ is the best predictor for $\theta_c(T)$, based on information set $\mathcal{F}_t$.

5.1.1 Comparison of the Mean Square Prediction Errors

Mathematically, MSPE($\tilde{P}_c(t)$) can be larger or smaller than MSPE($V_c(t)$). This is important. It says that when the market is efficient, the market forecast can be better or worse than the poll forecast. The argument that markets outperform polls because they are efficient is incorrect in general.

The results from the major polls are highly publicized, however, warranting the assumption that these results are included in the information set used by the market participants:

A1: $\mathcal{F}_t$ includes the poll results.

Under A1, because $V_c(t)$ is in the information set $\mathcal{F}_t$ and $\tilde{P}_c(t) = E(\theta_c(T) \mid \mathcal{F}_t)$,

$$MSPE(V_c(t)) = E(V_c(t) - \tilde{P}_c(t) + \tilde{P}_c(t) - \theta_c(T))^2 = MSPE(\tilde{P}_c(t)) + E(V_c(t) - \tilde{P}_c(t))^2. \tag{6}$$

That is, for every time $t < T$ and for any size poll, the equilibrium price $\tilde{P}_c(t)$ is better than the poll forecast $V_c(t)$. Although this is a strong result, a much stronger result holds. Letting $\mathcal{P}_t$ denote the information set containing the results from all the polls up to time t, it now follows again from A1 (for any measurable function f) that MSPE($\tilde{P}_c(t)$) $\leq$ $MSPE(f(\mathcal{P}_t))$. That is, the equilibrium price is better than any forecast obtained by combining information from polls at different times and/or different organizations.

It is also worth noting that the results above hold whether or not there is systematic measurement error in the polls, that is, whether or not assumptions S1–S5 hold. In other words, when each methodology works ideally, the market forecast is superior.

5.1.2 Comparison of the Bias

We know that $E(\tilde{P}_c(t) - \theta_c(T)) = 0$; that is, the prediction bias of the equilibrium price is zero. However, by Eq. (4), the prediction bias of $V_c(t)$, i.e., $E(V_c(t) - \theta_c(T)) = E[B(\theta_c(t))] + E[\theta_c(t) - \theta_c(T)]$, is not in general zero. If the polls work ideally, i.e., S1–S5 hold, the prediction bias of $V_c(t)$ reduces to $E[\theta_c(t) - \theta_c(T)]$, the expected value of the projection error. This result supports the potential utility of using the market time series to study electoral campaigns, as suggestd in Shaw and Roberts (2000). Note that assumption A1 was not used to obtain this result.

5.1.3 Comparison of the Variances

Although we have shown that $MSPE(\tilde{P}_c(t))$ is smaller than $MSPE(V_c(t))$ when A1 holds, the variance of $\tilde{P}_c(t)$ could nevertheless exceed that of $V_c(t)$. However, if assumptions S1–S5 hold and we also assume only (without using A1) that the electorate's preferences are less volatile on election day than at any other time $t < T$, that is,

A2: $Var(\theta_c(T)) \leq Var(\theta_c(t))$, for all $t < T$,

then it follows that

$$Var(\tilde{P}_c(t)) \leq Var(V_c(t)), \tag{7}$$

no matter how large the sample size. Note also that A2 does not imply that $Var(\theta_c(t))$ must decrease as t increases. The result, Eq. (7), may provide some explanation for the empirical finding that markets seem to exhibit less volatility than polls.

To show Eq. (7), note that with $B(\theta_c(t)) = 0$, $V_c(t) = \theta_c(t) + \varepsilon_c(t)$; thus

$$\begin{aligned} Var(V_c(t)) &= Var(\theta_c(t)) + Var(\varepsilon_c(t)) + 2Cov(\theta_c(t), \varepsilon_c(t)) \\ &= Var(\theta_c(t)) + Var(\varepsilon_c(t)), \end{aligned}$$

as $E[\varepsilon_c(t) \mid \theta_c(t)] = 0$. Therefore, by A2,

$$Var(\tilde{P}_c(t)) = Var(E[\theta_c(T) \mid \mathcal{F}_t]) \leq Var(\theta_c(T)) \leq Var(\theta_c(t)) \leq Var(V_c(t)),$$

regardless of the sample size.

5.1.4 Mean Square Prediction Error of $V_c(t)$

We have compared $MSPE(V_c(t))$ with $MSPE(\tilde{P}_c(t))$. We now consider the behavior of $MSPE(V_c(t))$ as t approaches T, i.e., as election day nears. We assume throughout that S1–S5 hold, implying $B(\theta_c(t)) = 0$; it then follows from Eq. (4) that

$$MSPE(V_c(t)) = MSPE(\theta_c(t)) + Var(\varepsilon_c(t)) + 2Cov(\theta_c(T) - \theta_c(t), \varepsilon_c(t)). \quad (8)$$

We showed above that $\varepsilon_c(t)$ is uncorrelated with $\theta_c(t)$. It would be odd were the sampling error from the poll at time t ($\varepsilon_c(t)$) correlated with the true vote at time T $\theta_c T$. Hence we assume

A3: $Cov(\varepsilon_c(t), \theta_c(T)) = 0$.

Consequently, Eq. (8) now reduces to

$$MSPE(V_c(t)) = MSPE(\theta_c(t)) + Var(\varepsilon_c(t)); \quad (9)$$

that is, if A3 holds, MSPE($V_c(t)$) is the variance of $\varepsilon_c(t)$ plus the MSPE due to projecting $\theta_c(t)$ forward to time T. Note that by taking a large sample $Var(\varepsilon_c(t))$ can be made arbitrarily small.

As election day nears, it is typically reasonable to assume that the electorate's preferences lock in, that is, $\theta_c(t)$ approaches $\theta_c(T)$. Consistent with this idea, Gelman and King (1993) find that $V_c(t)$ becomes close to $\theta_c(T)$ as t gets closer to T. There are a number of ways to formalize this idea. Here we assume:

A4: The sample path of the stochastic process $\theta_c(t)$ is continuous in time.

Because $|\theta_c(t)| < 1$, A4 implies

$$\lim_{t \uparrow T} MSPE(\theta_c(t)) = 0. \quad (10)$$

Thus, combining Eq. (9), which relies solely on A4, with Eq. (8), which relies on S1−S5 and A3, for t close to T, $MSPE(V_c(t)) \approx Var(\varepsilon_c(t))$. Note, however, because we do not know when t is close to T, it is necessary to make a subjective judgment (which could prove wrong) that this approximation is appropriate. In some elections, this might be the case many weeks before election day; in others, perhaps only a few days.

5.2 *Practice*

The mathematical results above were obtained assuming that M1–M6 and/or S1–S5 hold. In practice this will not be strictly true. However, the theoretical results provide a baseline for systematically discussing how the failure of one or more assumptions is likely to affect the relative performance of polls and markets. Thus we now briefly consider those assumptions most likely to fail and their impact on these results. Substantive considerations lead to the conclusion that even when the assumptions do not hold, the results obtained above will often hold approximately.

We first examine our assumptions about the market M1–M6. Suppose shares trade at the unique price $P_c(t) \in \mathcal{F}_t$. If expectations are heterogeneous, the market is not liquid, or some of the other market assumptions fail to hold, $P_c(t) \neq \tilde{P}_c(t)$ and

$$MSPE(P_c(t)) = MSPE(\tilde{P}_c(t)) + E(P_c(t) - \tilde{P}_c(t))^2. \quad (11)$$

Perhaps the most important case is when M3 fails to hold, that is, market participants have heterogeneous expectations. Nevertheless, aggregate expectations may be rational, in which case $P_c(t) = \tilde{P}_c(t)$. This can occur if marginal traders (see Forsythe et al. 1992) whose actions and subjective expectations are rational drive the price $P_c(t)$ toward $E(\theta_c(T) \mid \mathcal{F}_t)$, suggesting that in practice, $E(P_c(t) - \tilde{P}_c(t))^2$ may be small, at least in large markets. In other words, a smart subset of the participants will capitalize on the incorrect expectations of others by trading accordingly, driving the market price $P_c(t)$ toward $\tilde{P}_c(t)$.

In polling, while deviations from probabilistic sampling sometimes occur at the household stage, the potential bias due to nonresponse and the sytematic error introduced by polling organizations in allocating undecided voters and determining who is likely to vote (Erikson 1993; Voss et al. 1995) are of greater concern. Assumptions S2–S4, introduced to handle these sources of systematic measurement error, are not entirely plausible. However, even when these assumptions are incorrect, if nonresponse rates were low, the vast majority of eligible voters voted, and the percentage undecided were small, the bias would also be small. But with refusal rates of over 30% common in telephone polls (Crespi 1988), a substantial fraction of registered (eligible) voters not voting, and proportions undecided often exceeding 10%, the potential systematic measurement error is substantial. Much more work on modeling and correcting for this type of error is needed.

When systematic measurement error is present, $MSPE(V_c(t))$ is not in general a simple sum of nonnegative components, as in Eq. (9). Under assumption A3, Eq. (4) implies

$$\begin{aligned} MSPE(V_c(t)) = {} & E(\theta_c(T) - \theta_c(t))^2 + E[B(\theta_c(t))^2] \\ & + Var(\varepsilon_c(t)) - 2E[(\theta_c(T) - \theta_c(t))B(\theta_c(t))]. \end{aligned} \tag{12}$$

Using assumption A4, Eq. (12) reduces to

$$MSPE(V_c(t)) \approx Var(\varepsilon_c(t)) + E(B(\theta_c(t))^2), \quad \text{if } t \approx T. \tag{13}$$

In general, the term due to systematic measurement error should decrease as election day nears. First, the need for timely responses generally results in higher rates of nonresponse in the tracking polls conducted shortly before the election, increasing potential systematic measurement error. However, as nonrespondents are also less likely to vote, their exclusion can actually decrease systematic measurement error. Second, the proportion undecided may decrease as election day approaches, thereby decreasing the potential impact of allocation errors on forecasts from poll data; for example, in the Gallup/CNN/USA Today polls previously discussed, the percentage undecided declines to 3% in the last few polls. Further, those who are undecided close to the election are less likely to vote, so the commonly used scheme of allocating the percentage undecided proportionally to the vote shares of the major candidates (as we also did earlier) often works reasonably well.

The arguments above suggest that even when one or more of the assumptions justifying the market and the poll forecasts fail to hold, the ideal case results may still approximately hold, especially as the election approaches. Thus, also assuming A1, our theoretical analysis of the markets and polls continues to provide some explanation for the (often observed) superiority of the market forecast.

6 Discussion and Future Agenda

We have argued that election markets merit the attention of political scientists on both theoretical and practical grounds. We have rigorously developed, starting from first principles, the rationale for using market prices to forecast electoral outcomes. We have also developed the rationale for using poll-based forecasts. We have proposed the use of the mean square prediction error for comparing the predictive power of alternative forecasts. We have proved that the market forecast does not dominate poll-based forecasts in general. We have proved that the market forecast dominates poll-based forecasts if the polls are included in the information set of the market participants. Not only have we established this result for a given poll; the result holds for any (measurable) function of any set of polls. That is, the market outperforms even the best forecaster who bases his/her prediction on the polls. We have also considered the practical case in which one or more of the assumptions M1–M6 and/or S1–S5 fail to hold, arguing on substantive grounds that our theoretical results help to account for the empirical fact that markets often predict electoral results better than polls. We have also proved that the expected prediction error of the market forecast is 0 (not true for the polls) and that (under mild assumptions) the market forecast has smaller variance than a poll-based forecast. These results lend weight to the suggestion (Shaw and Roberts 2000) that the series of market prices is useful for studying campaign processes. Finally, articulating the types of assumptions that must hold in order for election markets to provide useful forecasts of electoral outcomes makes clear that it is these very same types of assumptions that more generally offer a theoretical foundation for the use of experimental markets to predict outcomes and aggregate information.

We are certainly not advocating that polling be abandoned. Polls are used for purposes other than forecasting, such as studying the relationship between voter characteristics and electoral preferences. Further, although a few authors have suggested otherwise (for example, Bohm and Sonnegård 1999; Forsythe et al. 1992), we believe, as per assumption (A1), that the polls are an important source of information used by market participants. We also note that in the studies just cited, restrictive assumptions about the form of the relationship between the market data and the polling data were made, and statistical tests based on the normal distribution were (inappropriately) applied to percentages.

The comparative framework we have constructed can be extended in several directions. Our results are readily applied to other payoffs dependent on the vote shares. Recalling that our results on IEM vote share markets are readily modified to apply to other IEM markets, such as the winner-take-all election markets, in which the payoff is \$1 per share of the candidate with the largest share of the popular vote, 0 otherwise, our framework can be used to compare winner-take-all forecasts of winners with forecasts obtained by other methods, for example, forecasting winners within each state and aggregating the results, as in Gelman and King (1993). Our framework can also be used to compare the performance of the market forecast with that of the longer-term forecasting methods mentioned in the introduction. Because a full treatment of this topic merits a separate paper, we do not discuss this further here.

We conclude by pointing out some limitations of our work. In deriving the equilibrium market prices, we assumed that participants used a common information set and held homogeneous expectations. Assumptions such as these, although common in finance and economics, most likely do not constitute an accurate description of the market participants. Clearly, the theoretical justification for using election markets would be strengthened were these assumptions removed. In that vein, perhaps the assumption of homogeneous expectations could be replaced by the assumption that participants hold "concordant" beliefs (Milgrom and Stokey 1982).

We have focused on theoretical issues in this paper. We have shown that the equilibrium market price is a better predictor than a public opinion poll (or a prediction made from combining the results of multiple polls) under the reasonable assumption that market participants know the poll results. Although we have also given reasons to expect actual market prices to predict better than poll-based forecasts, the case for using the market price as a predictor is clearly strongest when the market is in equilibrium. Thus it is useful to be able to ascertain whether or not the market is in this state. In future work, we intend to show how this issue can be empirically addressed, using the fact that the series of equilibrium prices constitutes a martingale sequence.

References

Beck, Nathaniel. 1992. "We Should be Modest." *Political Methodologist* 5:19–24.

Beck, Nathaniel. 2000. "Evaluating Forecasts and Forecasting Models of the 1996 Presidential Election." In *Before the Vote: Forecasting American National Elections.* eds. James E. Campbell and James C. Garand. Thousand Oaks, CA: Sage, pp. 161–167.

Berg, Joyce, Robert Forsythe, and Thomas Rietz. 1997. "What Makes Markets Predict Well? Evidence from the Iowa Electronic Markets." In *Understanding Strategic Interaction: Essays in Honor of Reinhard Selten,* ed. Wilfried Guth. Berlin: Springer-Verlag, pp. 444–463.

Blomberg, S. Brock, and Gregory D. Hess. 1997. "Politics and Exchange Rate Forecasts." *Journal of International Economics* 43:189–205.

Bohm, Peter, and Joakim Sonnegård. 1999. "Political Stock Markets and Unreliable Polls." *Scandinavian Journal of Economics* 101:205–222.

Brander, James A. 1991. "Election Polls, Free Trade, and the Stock Market: Evidence from the 1988 Canadian General Election." *Canadian Journal of Economics* 24:827–843.

Campbell, James E. 1992. "Forecasting the Presidential Vote in the States." *American Journal of Political Science* 36:386–407.

Campbell, James E., and James C. Garand, eds. 2000. *Before the Vote: Forecasting American National Elections.* Thousand Oaks, CA: Sage.

Chan, Nicholas T., Ely Dahan, Andrew W. Lo, and Tomaso Poggio. 2001. "Experimental Markets for Product Concepts." CBCL Memo 200, MIT Sloan School of Management.

Crespi, Irving. 1988. *Pre-election Polling: Sources of Accuracy and Error.* New York: Russell Sage Foundation.

Davis, Mark. 1997. "Option Pricing in Incomplete Markets." In *Mathematics of Derivative Securities,* eds. Michael A. H. Dempster and Stanley R. Pliska. Cambridge: Cambridge University Press, pp. 216–226.

Erikson, Robert S. 1989. "Economic Conditions and the Presidential Vote." *American Political Science Review* 83:567–573.

Erikson, Robert S. 1993. "Counting Likely Voters in Gallup's Tracking Poll." *Public Perspective,* March: 22–23.

Erikson, Robert S., and Christopher Wlezien. 1994. "Forecasting the Presidential Vote, 1992." *Political Methodologist* 5:10–11.

Fair, Ray C. 1978. "The Effect of Economic Events on Votes for President." *Review of Economics and Statistics* 60:59–172.

Fair, Ray C. 1996. "Econometrics and Presidential Elections." *Journal of Economic Perspectives* 10:89–102.

Forsythe, Robert, Murray Frank, Vasu Krishnamurthy, and Thomas W. Ross. 1998. "Markets as Predictors of Election Outcomes: Campaign Events and Judgement Bias in the 1993 UBC Election Stock Market." *Canadian Public Policy* 24:329–351.

Forsythe, Robert, Forrest Nelson, George R. Neumann, and Jack Wright. 1992. "Anatomy of an Experimental Political Stock Market." *American Economic Review* 82:1142–1161.

Forsythe, Robert, Thomas A. Rietz, and Thomas W. Ross. 1999. "Wishes, Expectations and Actions: A Survey on Price Formation in Election Stock Markets." *Journal of Economic Behavior and Organization* 39:83–110.

Gelman, Andrew, and Gary King. 1993. "Why Are American Presidential Election Campaign Polls So Variable When Votes Are So Predictable?" *British Journal of Political Science* 23:409–451.

Herron, Michael C., James Lavin, Donald Cram, and Jay Silver. 1999. "Measurement of Political Effects in the United States Economy: A Study of the 1992 Presidential Election." *Economics and Politics* 11:51–81.

Hull, John. 2000. *Options, Futures, and Other Derivative Securities,* 4th ed. Englewood Cliffs, NJ: Prentice Hall.

Karatzas, Ioannis, and S. G. Kou. 1996. "On Pricing of Contingent Claims under Constraints." *Annals of Applied Probability* 6:321–369.

Knight, Brian. 2003. "Are Policy Platforms Capitalized into Equity Prices? Evidence from the Bush/Gore 2000 Presidential Election." Unpublished manuscript, Brown University.

Kou, S. G., and Michael E. Sobel. 2001. "Hedging Electoral Risk." *Risk* 14:95–98.

Lau, Richard R. 1994. "An Analysis of the Accuracy of 'Trial Heat' Polls during the 1992 Presidential Election." *Public Opinion Quarterly* 58:2–20.

Lewis-Beck, Michael S., and Mary Stegmaier. 2000. "Economic Determinants of Electoral Outcomes." *Annual Review of Political Science* 3:183–219.

Little, Roderick J. A., and Donald B. Rubin. 1987. *Statistical Analysis with Missing Data.* New York: Wiley.

Lucas, Robert. 1978. "Asset Prices in an Exchange Economy." *Econometrica* 46:1429–1445.

Milgrom, Paul, and Nancy Stokey. 1982. "Information, Trade and Common Knowledge." *Journal of Economic Theory* 26:17–27.

Mitofsky, Warren J. 1998. "Was 1996 a Worse Year for the Polls than 1948?" *Public Opinion Quarterly* 62:230–249.

Mosteller, Frederick, Herbert Hyman, Phillip J. McCarthy, Eli S. Marks, and David B. Truman. 1949. *The Pre-election Polls of 1948.* New York: Social Science Research Council.

Newport, Frank. 1997. "The Pre-election Polls Performed Well in '96." *Public Perspective*, January: 50–51.

Panagakis, Nick. 1997. "Incumbent Races: A National Perspective." *Public Perspective*, January: 21–23.

Pennock, David M., Steve Lawrence, Finn A. Nielsen, and C. Lee Giles. 2001. "Extracting Collective Probabilistic Forecasts from Web Games." In *Proceedings of the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 174–183.

Rhode, Paul W., and Koleman S. Strumpf. 2003. "Historical Presidential Betting Markets." Unpublished manuscript, University of North Carolina at Chapel Hill.

Riley, William B., and William A. Luksetich. 1980. "The Market Prefers Republicans: Myth or Reality." *Journal of Financial and Quantitative Analysis* 15:541–560.

Rosenstone, Steven J. 1983. *Forecasting Presidential Elections.* New Haven, CT: Yale University Press.

Shaw, Daron R., and Brian E. Roberts. 2000. "Campaign Events, the Media and the Prospects of Victory: The 1992 and 1996 U.S. Presidential Elections." *British Journal of Political Science* 30:259–290.

Spann, Martin, and Bernd Skiera. 2003. "Internet-Based Virtual Stock Markets for Business Forecasting." *Management Science* 49:1310–1326.

Stokey, Nancy, and Robert Lucas. 1993. *Recursive Methods in Economic Dynamics.* Cambridge, MA: Harvard University Press.

Traugott, Michael W., and Paul J. Lavrakas. 2000. *The Voter's Guide to Election Polls*, 2nd ed. New York: Seven Bridges.

Voss, D. Stephen, Andrew Gelman, and Gary King. 1995. "Preelection Survey Methodology: Details from Eight Polling Organizations, 1988 and 1992." *Public Opinion Quarterly* 59:98–132.

Wlezien, Christopher, and Robert S. Erikson. 2000. "Temporal Horizons and Presidential Election Forecasts." In *Before the Vote: Forecasting American National Elections,* eds. James E. Campbell and James C. Garand. Thousand Oaks, CA: Sage, pp. 103–118.