Memoirs of an indifferent trader: Estimating forecast distributions from prediction markets

Joyce E. Berg
Department of Accounting, University of Iowa

John Geweke
Centre for the Study of Choice, University of Technology Sydney and
Department of Finance, University of Colorado

Thomas A. Rietz
Department of Finance, University of Iowa

Prediction markets for future events are increasingly common and they often trade several contracts for the same event. This paper considers the distribution of a normative risk-neutral trader who, given any portfolio of contracts traded on the event, would choose not to reallocate that portfolio of contracts even if transactions costs were zero. Because common parametric distributions can conflict with observed prediction market prices, the distribution is given a nonparametric representation together with a prior distribution favoring smooth and concentrated distributions. Posterior modal distributions are found for popular vote shares of the U.S. presidential candidates in the 100 days leading up to the elections of 1992, 1996, 2000, and 2004, using bid and ask prices on multiple contracts from the Iowa Electronic Markets. On some days, the distributions are multimodal or substantially asymmetric. The derived distributions are more concentrated than the historical distribution of popular vote shares in presidential elections, but do not tend to become more concentrated as time to elections diminishes.

Keywords. Forecasting, information systems analysis and design, probability distributions, Bayesian estimation, Iowa Electronic Markets.

JEL classification. C11, C93, D8, G1.

1. Introduction

A prediction market executes trades of contracts whose final values will be determined by the outcome of a specific future event. While the techniques developed here can be applied widely in prediction markets, the focus of this paper is the Iowa Electronic Markets (IEM) for U.S. presidential elections. Final contract values in an IEM presidential election market are determined by the popular vote share. For example, the final value

of an IEM vote-share (VS) contract is proportional to the indicated candidate's share of the popular vote for the major parties, while an IEM winner-takes-all (WTA) contract has value 1 if the indicated candidate wins a plurality of the popular vote and 0 otherwise. Other contracts with final values that depend on vote shares have also been traded in the IEM presidential election markets, as detailed in Section 2.

Prior to the determination of the final values of the contracts traded in a prediction market, their prices fluctuate as the future event approaches. For an idealized risk-neutral trader who has no transactions costs, the optimal allocation of a prediction market portfolio is directly related to contract prices and that trader's subjective probability distribution of the future event that determines final contract values. This trader's position in an IEM VS contract is determined by prices and her subjective expectation of the proportion of the major party popular vote won by the indicated candidate. Her position in an IEM WTA contract is determined by prices and her subjective probability that the indicated candidate receives the most votes.

The focus of our analysis is the subjective probability distribution of such an idealized trader who, moreover, endowed with *any* portfolio of IEM presidential election market contracts, would not reallocate that portfolio. Since, at any given time, traders observe the highest bid and the lowest ask price of VS contracts, these bid and ask prices bound the mean of this trader's distribution of the proportion of the popular vote going to each of the two major party candidates. Similarly, the WTA bid and ask prices bound this trader's subjective probabilities that each candidate's share will exceed one-half. We call this hypothetical individual the *indifferent trader*.

The objective of this paper is to learn from observed VS, WTA, and other contract prices what the subjective probability distribution of an indifferent trader would be were such a trader to exist. Our methods and results in no way rely on whether or not an indifferent trader actually exists. Neither do they depend on the distribution of beliefs among actual traders, an interesting but distinct issue (Manski (2006), Wolfers and Zitewitz (2006)).

Mapping contract prices into an indifferent trader's subjective probability distribution turns out to be a well framed but unexplored question with interesting answers. To begin, it immediately raises the questions of whether there exists such a probability distribution and, if so, whether the distribution is unique. Section 3 of the paper addresses the question of existence in a simple prediction market with only VS and WTA contracts, each with coincident bid and ask prices. Proposition 1 shows that existence is equivalent to a simple restriction on VS and WTA prices. The restriction is weak and easily satisfied by the IEM presidential election market prices studied here.

It would be convenient to assume that the indifferent trader's distribution has a simple parametric form.[1] Section 3 shows that common parametric forms (beta and logistic normal) impose much stronger restrictions on combinations of VS and WTA prices.

---

[1]Often, researchers estimate forecast distributions from prediction markets by imposing such forms. For example, Berg, Neumann, and Rietz (2009) assumed a log normal distribution for market capitalization after the Google initial public offering. Gruca, Berg, and Cipriano (2008) used a normal distribution for movie box office revenues. Berg, Nelson, and Rietz (2008) assumed a logistic normal distribution for U.S. presidential vote shares, while Page (2008) assumed a normal distribution as did Leigh and Wolfers (2006) implicitly.

These restrictions are violated in much of our data and in the data used in other studies as well.[2] Furthermore, in Section 3, we show that IEM presidential election market VS and WTA prices rule out symmetric distributions on many days.

These restrictions on probability distributions imposed by prediction market contract prices—which, to the best of our knowledge, are all new results—led us to a nonparametric representation of the indifferent trader's probability distribution of the major party shares of the popular vote. Specifically, this distribution is cast as discrete on the points $a_i = (i - 0.5)/n$ $(i = 1, \ldots, n)$. In a nonparametric environment there are, in general, many distributions that are consistent with bid and ask prices as long as the conditions of Proposition 1 are satisfied. These conditions are satisfied on all of the days studied subsequently in this paper.

Our approach, detailed in Section 4, is to order the set of possible subjective distributions using a prior distribution that favors probability distributions that are smooth and concentrated, the tension between these criteria being governed by the specific choice of the prior. Because we wish to produce these distributions repeatedly, reliably, and in real time, Section 4 focuses on a prior distribution for which finding the mode of the posterior distribution amounts to the solution of a quadratic programming problem in n variables.

Section 5 uses this approach to examine the probability distributions of the indifferent trader, corresponding to the posterior mode, at midnight on the 100 days leading up to the presidential elections of 1992, 1996, 2000, and 2004. We find that these distributions are often notably asymmetric, many (though a minority) of them are bimodal, and a few are multimodal. Significantly, neither the volatility nor the variance of these distributions changes systematically as election day approaches.

The main conclusions of our research, elaborated in the final section, are (i) common parametric distributions must be avoided when inferring distributions from contract prices in prediction markets, (ii) nonparametric probability distributions of the events priced by these contracts can be found rapidly and reliably in real time, (iii) the subjective probability distribution of the indifferent trader in the four presidential elections considered is substantially more concentrated than the distribution of realized vote shares in the elections of 1868–1988, and (iv) the evolution of distributions leading up to the elections is more consistent with inference about an evolving latent variable than with classic learning about a fixed but unknown parameter.

2. The IEM presidential election markets

The University of Iowa runs the Iowa Electronic Markets (IEM)—real-money futures markets—through the Internet for teaching and research purposes. Traders worldwide can establish accounts with initial investments between \$5 and \$500. Since the IEM is described in detail at its website (http://www.tippie.uiowa.edu/iem/) and elsewhere

---

[2]This is why Berg, Nelson, and Rietz (2008) frequently cannot estimate implied volatilities for the logistic normal distribution they use. Figure 2 in Berg, Neumann, and Rietz (2009) sometimes shows evidence of a multimodal distribution. Distributions in Chen and Plott (2002) show frequent asymmetries and, sometimes, multiple modes.

(Forsythe, Nelson, Neumann, and Wright (1992), Forsythe, Rietz, and Ross (1999), Berg, Forsythe, and Rietz (1997), Oliven and Rietz (2004)), we will be brief here.

IEM markets are organized as continuous electronic double auctions. At any time, traders can place bids to buy contracts or asks to sell contracts into price and time ordered queues. At any given time, standing best (highest) bids and best (lowest) asks represent the best, immediately available prices that no trader is currently willing to accept. Traders who wish to trade immediately can accept outstanding best bids or asks and, as a result, execute a trade. This reveals the next best bid or ask. Finally, traders can buy or sell "unit portfolios" to or from the exchange for $1 at any time. Unit portfolios consist of one of each contract in a market. For example, in 2004 the vote-share market had two contracts, one for the Republican and one for the Democrat. By design, the sum of liquidation values of these contracts will always equal $1.

Prior to each of the four U.S. presidential elections 1992–2004, the IEM traded both vote-share (VS) and interval (INT) contracts.[3] Although the exact specification of the contracts changed from election to election, all give information about the popular vote shares of the candidates. The Appendix (Supplemental Material (Berg, Geweke, and Rietz (2010))) details the specific definitions of all contracts studied. It also sets up a common nomenclature for the contracts, which we will use here. Each market had several hundred active traders. Prospectuses for the markets are available from the IEM website.

After the election, VS contracts liquidate (pay their owners) $1 times the (appropriately defined) share of the popular vote cast for the associated candidate. For example, the contract we denote by $VS.D|DR$ paid $1 multiplied by the Democratic nominee's (i.e., John Kerry's) share of the two-party vote in 2004 and the contract we denote by $VS.R|DR$ paid $1 multiplied by the Republican nominee's (i.e., George Bush's) share. Since Kerry took 48.8% of the two-party vote to Bush's 51.2%, the $VS.D|DR$ contract paid $0.488 and the $VS.R|DR$ contract paid $0.512. The unit portfolio consisted of one of each contract and was worth $1.

Interval contracts generally pay $1 if the vote share falls in a given range and $0 otherwise. For example, a contract may pay $1 if the associated candidate takes the plurality of the vote, like the contract we denote by $W.D|DR$ in 2004 (which paid $1 if the Democrat (Kerry) took more popular votes than the Republican (Bush)). The unit portfolio consisted of this and a corresponding contract for the Republican party. We refer to these interval contracts as winner-takes-all (WTA) contracts. Sometimes interval contracts subdivided the vote ranges. For example, in 2004, the IEM subdivided the payoff ranges of the original $W.D|DR$ and $W.R|DR$, effectively splitting each into two new contracts. The contract we denote by $(50, 52).R|DR$ paid $1 if the Republican (Bush) took between 50% and 52% of the two-party vote in 2004 and $(52, 100).R|DR$ paid $1 if he took more than 52%. Similar contracts were defined for the Democrat (Kerry) taking 50% to 52% and more than 52%. These four contracts formed a new unit portfolio.

Taken together, highest bid and lowest ask quotations for vote-share and interval contracts restrict the indifferent trader's subjective distribution of vote shares. Vote-share contracts provide information about the mean of the distribution, while interval contracts provide information about the probability mass over specified ranges. For

---

[3]In 1988, the IEM traded only VS contracts.

each contract, we record the midnight highest outstanding bids and lowest outstanding asks on each of the last 100 days of the market.[4] Our indifferent trader's expectation of f must lie between the bid and ask prices of the vote-share contracts, and similarly this trader's subjective probabilities of each event in the interval markets must lie between the bid and ask of the associated contract.

3. Multiple contracts and probability distributions

We begin by investigating implied distribution restrictions in a simple prediction market for an event that is a fraction $f \in [0, 1]$. The market studied is simple, with just two contracts: the liquidation value of one contract is the realized value of f; the liquidation value of the other contract is 1 if $f > 0.5$ and 0 otherwise. In this section, we ignore complications due to bid–ask spreads and assume that the contracts have unique prices v (for vote share) and w (for winner take all), respectively. The Appendix (Supplemental Material (Berg, Geweke, and Rietz (2010))) contains proofs of all of the propositions.

The hypothesis that, for some random variable f, it is simultaneously the case that $v = E(f)$ and $w = P(f > 0.5)$ places restrictions on v and w.

Proposition 1. *The set S of possible* (v, w) *is*

$$S = \{(v, w) : 0 < v < 1; 0 \leq w \leq 1; 2v - 1 < w < 2v\}.$$

The set S of Proposition 1 is the parallelogram outlined by the heavy lines in Figure 1(a). The points in this figure are the 399 combinations (v, w) for the days studied. The latter are computed as means of bounds whose construction is detailed at the start of the next section. Clearly all of the (v, w) values are well within S and, hence, satisfy these minimal essential restrictions.

Proposition 2. *Suppose that the distribution of f is symmetric. The set S of possible* (v, w) *is* $S_s = S_1 \cup S_2 \cup S_3 \cup S_4 \cup S_5$, *where*

$$S_1 = \{(v, w) : 1/4 < v < 1/2, 0 < w < 1/2\},$$

$$S_2 = \{(v, w) : 1/2 < v < 3/4, 1/2 < w < 1\},$$

$$S_3 = \{v = w = 1/2\},$$

$$S_4 = \{w = 0, v < 1/2\},$$

$$S_5 = \{w = 1, v > 1/2\}.$$

---

[4]While the market never actually closes, the IEM commonly reports the last trade prices before midnight as the "closing" prices for the day. Midnight bids and asks incorporate the accumulated information from the day. While it can be changed by traders, bids and asks expire by default at midnight one or two days after they are placed (depending on the election year). This results in a pattern of narrowing bid–ask spreads during a day up to midnight. After midnight, they may widen. In 1992, data for the 65th day, September 28, were lost. This accounts for the gap shown in many of our graphs. There are 99 days of contract prices for 1992 and 399 for all four elections together. Attesting to the liquidity of these markets, for most contracts traded since 1996, the bid–ask spread is less than $\frac{1}{2}$ cent. For some of these contracts, most spreads are 1 mil—the lowest possible since IEM trading prices are integer multiples of a mil.
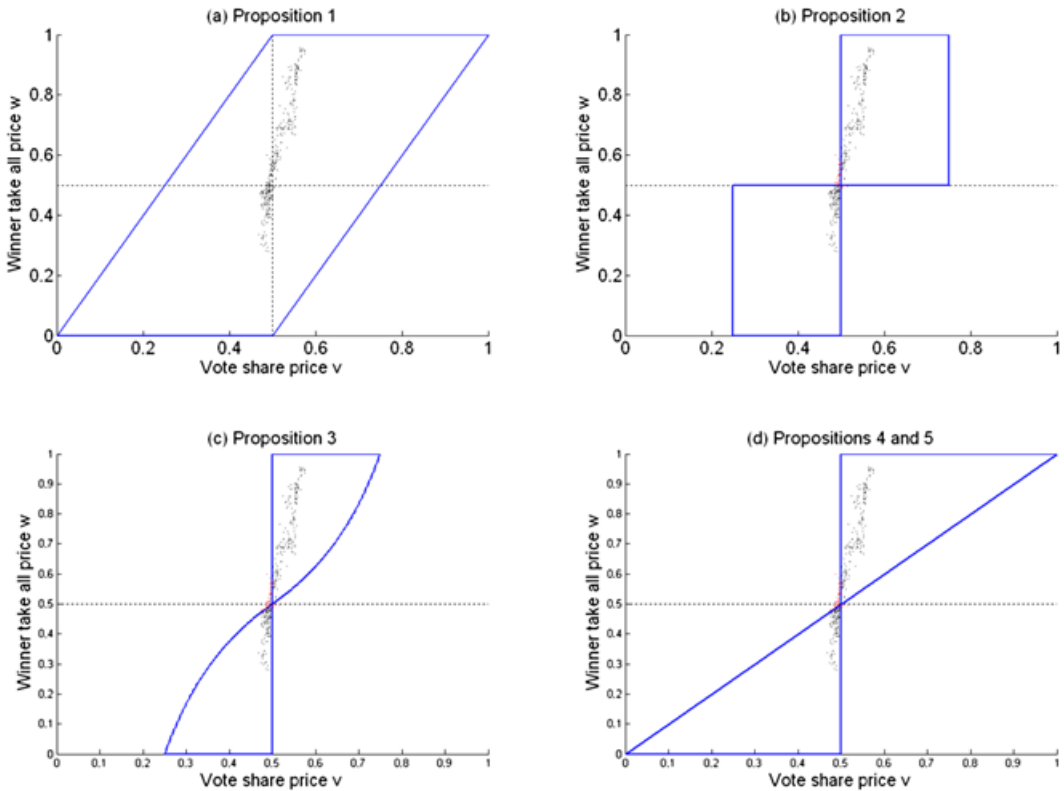
FIGURE 1. In each panel the interior of the region defined by the heavy lines indicates (v, w) combinations consistent with a distributional assumption. The dots indicate the midpoints of contract bid and ask prices for VS and WTA contracts before the four elections.

The sets S_1 and S_2 of Proposition 2 are the two rectangles each outlined by the heavy lines in Figure 1(b). Of the 399 combinations (v, w) in the sample, 44 (11.0%) are not in S_s—combinations in which one of the two prices was less than 0.5 while the other price exceeded 0.5. Thus for a significant part of our sample, the center of the bid–ask range is inconsistent with a symmetric distribution of f; for some of these sample points, the entire bid–ask range lies outside of S. The strongest violation of the constraints of Proposition 2 occurs on October 2, 2000, when $v = 0.4875$ and $w = 0.6005$. Section 5 returns to some of these cases in more detail.

The set of symmetric unimodal distributions places even more restrictions on the set of corresponding (v, w).

PROPOSITION 3. *Suppose that f is a continuous random variable with a symmetric distribution that has a global mode at $v = E(f)$. The set S of possible (v, w) is given by $S_u = S_{u1} \cup S_{u2} \cup S_3$, where*

$$S_{u1} = \{(v, w) : 1/4 < v < 1/2, 0 < w < 1 - (4v)^{-1}\},$$

$$S_{u2} = \{(v, w) : 1/2 < v < 3/4, [4(1 - v)]^{-1} < w < 1\},$$

and S_3 is defined in Proposition 2.

The set S_u of Proposition 3 is outlined by the heavy lines in Figure 1(c). In addition to the points excluded by Proposition 2, this result also excludes points (v, w) both less than 0.5 or both exceeding 0.5, but for which w is too close to 0.5. There are 62 points (15.5% of the total 399 days) in Figure 1(c) that violate the conditions of Proposition 3. The strongest violation of the constraints of Proposition 3 occurs on August 17, 2004, when $v = 0.5150$ and $w = 0.4948$; for consistency with Proposition 3, v would have to be reduced by 0.0150 or w would have to be increased by 0.0207.

Beta and logistic normal are two common distributions for random variables on the unit interval that need not be symmetric, but these distributions are also frequently inconsistent with the IEM data. The difficulties again involve restrictions placed on (v, w) near the point $(0.5, 0.5)$.

The beta distribution arises naturally when the information about a probability is updated with binary outcomes (Zellner (1971, Section 2.13)), for example, a poll of voting intentions.

PROPOSITION 4. *Suppose that f has a beta distribution. The set S_* of possible (v, w) is $S_* = S_{*1} U S_{*2} U S_3$, where*

$$S_{*1} = \{(v, w) : 0 < v < 1/2, 0 < w < v\},$$

$$S_{*2} = \{(v, w) : 1/2 < v < 1, v < w < 1\},$$

and S_3 is defined in Proposition 2.

The set S_* of Proposition 4 is outlined by the heavy lines in Figure 1(d). Notice that $S_u \subset S_* \subset S$. There are 59 days (14.8% of the total) in Figure 1(d) that are inconsistent with a beta distribution for f. The strongest violation of Proposition 4 is on August 17, 2004, the same date noted in connection with Proposition 3. For the beta distribution, however, an increase of 0.0202 in w would be required for consistency.

In previous work (e.g., Berg, Nelson, and Rietz (2008)), a log normal distribution has been used to analyze prediction markets. Here, a logistic normal distribution would be more appropriate.

PROPOSITION 5. *Suppose that f has a logistic normal distribution,* $\log[f/(1 - f)] \sim N(\mu, \sigma^2)$. *The set of possible (v, w) is S_* defined in Proposition* 4.

Proposition 1 and Figure 1(a) suggest that there are many distributions for f that will be consistent with the IEM data. Propositions 2 and 3 and Figure 1(b) and 1(c) show that the distributions cannot all be symmetric. Proposition 4 and Figure 1(d) show that many contract prices are inconsistent with beta distributions, and Proposition 5 and Figure 1(d) show inconsistencies with logistic normal distributions. These findings require that we permit more flexible distributions. The next section considers the class of all probability distributions for f, building on the ideas that a reasonable distribution should have a smooth probability density function and not be too dispersed.

4. A PRACTICAL FLEXIBLE MODEL FOR SUBJECTIVE DISTRIBUTIONS

We begin by modeling the Democratic fraction of the two-party vote f as a discrete random variable defined on a set of uniformly spaced points $a_i = (i - 0.5)/n$ $(i = 1, \ldots, n)$. While this cannot literally be true, by taking n sufficiently large, it is possible to approximate all relevant functions of the continuous random variable f. (All of the computations here take $n = 100$. None of the results reported subsequently changes substantively for larger values of n.) Denote the support of f by the $n \times 1$ vector $\mathbf{a} = (a_1, \ldots, a_n)'$. Then the distribution of f on day t is given by the corresponding $n \times 1$ vector $\mathbf{p}_t = (p_{t1}, \ldots, p_{tn})'$. For each election, $t = 1, \ldots, 100$ with $t = 100$ denoting the day before the election.

Our strategy is to cast the problem with a Gaussian prior distribution for $\mathbf{p}_t$ subject to the constraint that $\mathbf{p}_t$ is contained in the unit simplex of $\mathbb{R}^n$:

$$\mathbf{p}_t \in \mathcal{P}^0 = \left\{ \mathbf{p} : \mathbf{p} \geq \mathbf{0}, \sum_{i=1}^{n} p_{ti} = 1 \right\}.$$

Contract bid and ask prices provide additional inequality constraints as detailed in the Appendix. Intuitively, bids bound expectations from below, while asks bound them from above.[5] Table 1 shows the aspects of the distributions that are bounded by the traded contracts for all elections. We work with bounds on the Democratic fraction of the two-party vote in all four elections except 2000, in which it is the fraction of the three-party vote.[6]

The last column of Table 1 indicates the median bid–ask spread constructed in this way. Spreads in all of the 1996, 2000, and 2004 markets, and the 1992 VS market are a penny or less, and in some cases, they are only 1 mil—the smallest bid–ask spread possible. In relative terms, these spreads are comparable to those in markets for large-cap equities. Spreads in the 1992 INT markets are somewhat larger. There is no systematic tendency for spreads to increase or decrease as the election approaches.

Because of bid–ask spreads, restrictions on $\mathbf{p}_t$ come in natural pairs, for example, $VS.Db \leq E(f) \leq VS.Da$ for the vote-share contract (VS) that pays in proportion to the fraction of the popular vote received by the Democratic nominee (D), where b and a denote the highest bid and lowest ask, respectively. But, because contracts are part of a unit portfolio, there is another set of bounds. Traders can effectively "buy" a VS.D in two ways: (i) purchase it at the ask (VS.Da) or (ii) buy the unit portfolio for \$1 and sell the VS.R contract, giving a net price of $1 - VS.Rb$ to hold one additional share of VS.D. Traders can also "sell" a VS.D either directly at the bid or indirectly using a portfolio

---

[5]While the indifferent trader incurs no transactions costs and, hence, the bids and asks form the bounds, it is easy to adjust for transactions costs. One would simply lower the bound implied by the bid by subtracting the (implicit or explicit) transactions costs and raise the upper bound implied by the ask similarly. Increasing the bounds reduces the information content of the market and may lead to slightly smoother, more compact distributions, but otherwise leaves the procedure unchanged.

[6]The specific treatment of 2000 arises from the definitions of contracts for that election.

TABLE 1. Bounds on the vote-share distribution implied by each market.

Election	Contracts	Bounds Obtained From the Market	Median Spread Between Upper and Lower Bounds
1992	VS	$E(f)$	0.005
	WTA[*]	$P(f > 0.5)$	0.013
	INT1	$P(f > 0.46)$	0.013
	INT2[*]	$P(f > 0.50)$	0.013
	INT3	$P(f > 0.54)$	0.013
	INT4[**]	$P(f > 0.58)$	0.014
1996	VS	$E(f)$	0.006
	WTA	$P(f > 0.5)$	0.004
2000	VS	$E(f)$	0.003
	WTA	$P(f > 0.5)$	0.004
2004	VS	$E(f)$	0.003
	WTA[***]	$P(f > 0.5)$	0.010
	INT1[***]	$P(f < 0.48)$	0.001
	INT2[***]	$P(0.48 < f < 0.50)$	0.001
	INT3[***]	$P(0.50 < f < 0.52)$	0.002
	INT4[***]	$P(f > 0.52)$	0.001

Note: The asterisks indicate that in 1992, the IEM traded redundant contracts that each paid off based on the 50% vote-share cutoff. We computed the bounds using bids and asks from both markets. The spread reported here is the median smallest spread across all the bids and asks for all the contracts based on the 50% cutoff. The double asterisks denote that trading began 21 days before the election. The triple asterisks indicate the winner-takes-all contracts in 2004 based on a 50% vote-share cutoff were split into contracts associated with the $<48\%$, 48%–50%, 50%–52%, and $>52\%$ intervals. The split occurred 58 days before the election.

transaction. This leads to the second set of bounds: $1 - VS.Ra \le E(f) \le 1 - VS.Rb$. As our bounds, we use the more restrictive prices (lower buying price and higher selling price)

$$\underline{v}_t = \max(VS.Db, 1 - VS.Ra) \le E(f) \le \min(VS.Da, 1 - VS.Rb) = \overline{v}_t. \tag{1}$$

The corresponding restriction on the $n \times 1$ vector $\mathbf{p}_t$ in our model is

$$\mathbf{p}_t \in \mathcal{P}_t^1 = \{\mathbf{p}_t : \underline{v}_t \le \mathbf{a}'\mathbf{p}_t \le \overline{v}_t\}.$$

If the lower and upper bounds in (1) are the same, as occurred on a few dates in our sample, then the two inequality constraints are replaced by a single equality constraint.[7] Since the conventional quadratic programming problem (and software) incorporates equality as well as inequality constraints, this poses no difficulty. By extension, our procedures apply in markets with a market-maker rather than a continuous double auction. Finally, the second set of bounds is necessitated by the unit portfolio structure of the IEM. However, if unit portfolios did not exist, bounds could be derived from a single set of bids and asks.

---

[7]Occasionally, when a no-arbitrage restriction is violated, the highest lower bound exceeds the lowest upper bound. Oliven and Rietz (2004) showed that such arbitrage opportunities are quickly exploited. In these cases, we set both bounds to their midpoint, creating a single equality constraint.

The bounds for interval contracts are constructed according to the same principles, but, depending on the configuration of contracts offered and the number of candidates, the nature of the potential lower and upper bounds varies from year to year and case to case. In the same way the bid and ask prices of the WTA contract, which traded preceding each of the presidential elections we study, provide restrictions $\underline{w}_t \leq P(f \geq 1/2) \leq \overline{w}_t$ at the end of each trading day t. The corresponding restriction on $\mathbf{p}_t$ is $\mathbf{p}_t \in \mathcal{P}_t^2$ with $\mathcal{P}_t^2 := \{\mathbf{p}_t : \underline{w}_t \leq \sum_{i=n/2+1}^{n} p_{ti} \leq \overline{w}_t\}$. Other interval contracts, indicated in Table 1, provide similar restrictions. The restrictions on $\mathbf{p}_t$ from the bid and ask prices of all contracts at the end of day t can be expressed as

$$\mathbf{p}_t \in \mathcal{P}_t = \mathcal{P}^0 \cap \mathcal{P}_t^1 \cap \cdots \cap \mathcal{P}_t^{n_t},$$

where n_t is determined by the number of contracts traded. (For example, $n_t = 5$ in October 1992 and $n_t = 2$ in 2004. The Appendix details the construction of inequality constraints for each date from contract bid and ask prices.) The set $\mathcal{P}_t$ is the intersection of $\mathcal{P}^0$, which is the unit simplex in $\mathbb{R}^n$, with $2n_t$ additional inequality restrictions on linear functions of $\mathbf{p}_t$. Table 1 lists the properties of the distributions bounded by contract bids and asks preceding each election along with the median spreads between the bounds. The Appendix contains the derivations of the specified bounds for each year. As it shows, bounds can be obtained for a wide range of contract specifications within the IEM itself and the principles followed can be easily adapted to other prediction markets.

We restrict ourselves to a Gaussian prior distribution for the 100×1 parameter vector $\mathbf{p}_t$ for pragmatic reasons. If the prior distribution is Gaussian, then the posterior distribution is also Gaussian subject to the linear equality and inequality constraints just discussed. Since the logarithm of a Gaussian probability density is quadratic, finding the mode of the posterior distribution amounts to the solution of a quadratic programming problem. The solution can be computed in 0.4 seconds for the most complicated set of bounds in our data using standard desktop computing software and hardware, and it, therefore, can be implemented in real time in response to changing bid and ask prices— as we, in fact, did for the 2008 IEM presidential election contracts while the markets were running. The mode of the posterior distribution is our estimate $\widehat{\mathbf{p}}_t$ of the 100×1 parameter vector $\mathbf{p}_t$.

Our prior distribution reflects the belief that smoother, more concentrated distributions are more reasonable than rougher, more dispersed distributions. Therefore, the prior distribution has two independent components that reflect the beliefs that (i) p_{ti} is a smooth function of i and (ii) p_{ti} is likely to be smaller the larger is $|\mathbf{a}'\mathbf{p}_t - p_{ti}|$.

The first component of the prior distribution (smoothness) is the conventional visual smoothness prior of Shiller (1984) in which the first derivative of a continuous function follows a Wiener process with no drift and diffusion parameter τ^2: smaller values of τ^2 correspond to smoother functions. This implies a Gaussian prior for n regularly spaced points with mean zero and precision (inverse variance) matrix of the form $\tau^{-2}\mathbf{R}'\mathbf{GR}$. In this expression, $\mathbf{G}$ is an $(n-2) \times (n-2)$ tridiagonal matrix with $g_{ii} = 1/3$ ($i = 1, \ldots, n-2$) and $g_{i,i+1} = g_{i+1,i} = 1/6$ ($i = 1, \ldots, n-3$), and $\mathbf{R}$ is an $(n-2) \times n$ matrix with three nonzero elements $r_{ii} = 1$, $r_{i,i+1} = -2$, and $r_{i,i+2} = 1$ in each row i. In our application, $n = 100$ is the number of points in the unit interval used in the problem. Since the

interval between units has length $1/n$, that would call for scaling the precision matrix by n if $\mathbf{p}$ remained constant as n varied, but the constraints $\mathbf{p}_t \geq \mathbf{0}$ and $\sum_{i=1}^{100} p_{ti} = 1$ require that the elements of $\mathbf{p}_t$ are proportional to $1/n$ as well, for another factor of n^2. Therefore, the smoothness prior is Gaussian with precision matrix $\mathbf{H}_s(n, \tau^2) = n^3 \tau^{-2} \mathbf{R}'\mathbf{GR}$ and mean vector $\boldsymbol{\mu}_s = \mathbf{0}$. As explained in Geweke (2005, Section 5.4.1), this prior distribution for functional smoothness can be extended to any value of n and to a continuous function in a logically consistent manner. This property is unique to this prior distribution and provides a compelling reason for this particular form of the smoothness prior.

The second component of the prior distribution (concentration) severely penalizes substantial values of p_{ti} relatively far from the mean and has the kernel

$$\exp\left[-\sum_{i=1}^{n} \exp(c|a_i - \mathbf{a}'\mathbf{p}_t|) p_{ti} \right],$$

where c is a prior hyperparameter that regulates the size of the penalty. Since the log kernel is not quadratic in $\mathbf{p}_t$ this prior does not meet the objective of maximizing the log posterior by means of quadratic programming. We avoid this problem by substituting the mean v_t of the VS contract bounds for $\mathbf{a}'\mathbf{p}_t$. Given the small spreads in bounds for VS contracts, the effect of substituting v_t for $\mathbf{a}'\mathbf{p}_t$ should be small. The log kernel for this component of the prior distribution thus has the form $\mathbf{g}(c)'\mathbf{p}_t$, with $g_i(c) = -\exp(c|a_i - v_t|)$.

Incorporating both components, the log prior density kernel therefore has the form

$$Q(\mathbf{p}; \tau^2, c, \omega, n) = -\omega \frac{1}{2} \mathbf{p}_t' \mathbf{H}_s(n, \tau^2) \mathbf{p}_t - (1 - \omega) \exp[g(c)]' \mathbf{p}_t,$$

where $\omega \in (0, 1)$ is the weight on the smoothness component of the prior relative to concentration. Because $\text{rank}[\mathbf{H}_s(n, \tau^2)] = n - 2$, this is not a regular quadratic programming problem. To bring it into regular form, replace $\mathbf{H}_s(n, \tau^2)$ with the full rank matrix

$$\mathbf{H}(n, \tau^2) = \mathbf{H}_s(n, \tau^2) + n^{-1}[\mathbf{ee}' + \mathbf{aa}'],$$

where $\mathbf{e}$ is an $n \times 1$ vector of units. This changes the problem very little because $\mathbf{e}'\mathbf{p}_t = 1$ and $\mathbf{a}'\mathbf{p}_t$ is tightly constrained by the VS bids and asks.

The log likelihood function is $\mathcal{L}_t(\mathbf{p}_t) = 0$ for $\mathbf{p}_t \in \mathcal{P}_t$ and $\mathcal{L}_t(\mathbf{p}_t) = -\infty$ for $\mathbf{p}_t \notin \mathcal{P}_t$. Hence, the mode of the posterior density is

$$\widehat{\mathbf{p}}_t = \underset{\mathbf{p} \in \mathcal{P}_t}{\arg\max} \, Q(\mathbf{p}; \tau^2, c, \omega, n).$$

Since Q is a quadratic function and $\mathcal{P}_t$ is the intersection of linear inequality constraints, $\widehat{\mathbf{p}}_t$ is the solution of a conventional quadratic programming problem.

It remains to set the hyperparameters ω, τ^2, and c. Since the solution of the problem is invariant to scaling Q and changes in τ^2 scale $\mathbf{p}_t'\mathbf{H}\mathbf{p}_t$, ω and τ^2 are redundant. Without loss of generality, take $\omega = 1/2$. We experimented with the values of τ^2 and c, examining the effects of various combinations of these hyperparameters on the posterior

maximums for $\mathbf{p}_t$. We settled on the choices $\tau = 1$ and $c = 20$. The Appendix discusses sensitivity of the results to the selection of hyperparameters in more detail.

From a formal Bayesian perspective, $\widehat{\mathbf{p}}_t$ is the estimate of $\mathbf{p}_t$ given a zero–one loss function (Bernardo and Smith (1993, Section 5.1), Geweke (2005, Section 2.4)). From a formal sampling-theoretic perspective, the estimate is the maximum likelihood estimate of $\mathbf{p}_t$ using the penalized likelihood function

$$\mathcal{L}_t(\mathbf{p}_t) + Q(\mathbf{p}; \tau^2, c, \omega, n).$$

Under either the Bayesian or non-Bayesian interpretation, $\widehat{\mathbf{p}}_t$ is a point estimate of the indifferent trader's subjective probability distribution. Of course there are many other probability distributions $\mathbf{p}_t$ that are not excluded by contract bid and ask prices, and the point estimate conveys no sense of subjective uncertainty (in a Bayesian interpretation) or a sampling-theoretic distribution (in a non-Bayesian interpretation).

It is a straightforward but computationally time-intensive task to access the entire posterior distribution of $\mathbf{p}_t$ by means of a conventional Markov chain Monte Carlo algorithm for multivariate normal distributions subject to linear constraints (Geweke (1991)). Doing so with the speed and reliability required for real-time public reporting of $\mathbf{p}_t$, one of our primary objectives, raises challenges well beyond the scope of this paper.

The sampling-theoretic (non-Bayesian) distribution of $\widehat{\mathbf{p}}_t$ presents more fundamental problems. The relevant construct for asymptotic distribution theory is the introduction of successive contracts. For example, interval contracts with payment thresholds that become dense on the unit interval will render $\widehat{\mathbf{p}}_t$ superconsistent and, more generally, convergence rates will be sensitive to the structure of contracts. None of that is relevant here: the number of observations (i.e., inequality restrictions from bid and ask prices) is as few as 4 and no more than 12, while the number of parameters estimated is $n = 100$.

5. Subjective distributions of the indifferent trader

For the 100 days preceding the presidential elections from 1992 to 2004, we use daily closing bids and asks and the methods from Section 4 to compute $\widehat{\mathbf{p}}_t$. These are our estimates of the indifferent trader's distribution of the Democratic popular vote share, f. Figure 2 shows the results for the day preceding each election. The distribution is given by the solid curve in each panel. The solid vertical line segments at the top of each panel provide the lower and upper bounds of the VS contracts, and the dashed vertical line segments do the same for the WTA contracts.[8] The circle in the bound interval for the VS price is centered at the mean of the distribution (i.e., $\mathbf{a}'\widehat{\mathbf{p}}_t$). The square is centered at the probability that the Democrat receives more than half the popular vote in the election (i.e., $\sum_{i=51}^{100} \widehat{p}_{ti}$). Finally, for reference, the light diamonds plot out show a logistic normal distribution fit to the Democratic share of the two-party popular vote in the 1868–2004 presidential elections and centered at zero.

[8]The bounds for the VS contracts shown in Figures 2, 3, and 4 refer to the bounds on $E(f)$ calculated as indicated in the Appendix. The bounds for the WTA contracts in these figures refer to the bounds on $P(f > 0.5)$ calculated as given in the Appendix for 1992, 1996, 2000, and the first 58 days in 2004. For the remaining 42 days in 2004, the WTA bid is the sum of the lower bounds for $P(0.50 < f < 0.52)$ and $P(f \geq 0.52)$.
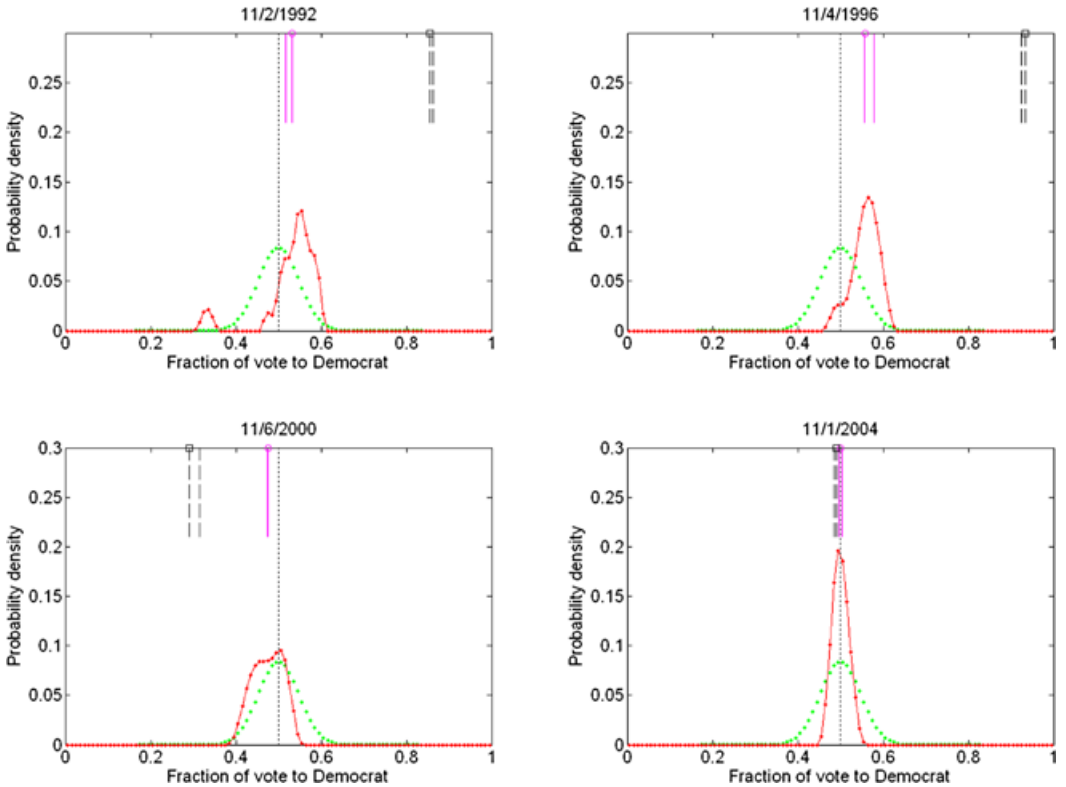
FIGURE 2. Estimated probability distribution at midnight preceding election day, with upper and lower bounds for vote-share and winner-takes-all contracts.

Using the same graphical presentation, we provide the distributions for all 399 days preceding the elections at http://www.tippie.uiowa.edu/iem/research/Bayesian/. Similar graphs were updated each minute during the 2008 presidential election market at the same website. Here we summarize some aspects of these distributions that are relevant for the issues motivating this study and for the use of our methods in monitoring prices in prediction markets more generally.

5.1 *Asymmetry and multiple modes*

Our procedure allows for asymmetric distributions. Section 3 shows that many of the combinations of VS and WTA prices observed are inconsistent with beta and logistic normal distributions, as well as with symmetric distributions. In particular, the combinations $(v < 0.5, w > 0.5)$ and $(v > 0.5, w < 0.5)$ are inconsistent with any of these distributions. For 23 of the 399 days, we observe these combinations not only for the mean of the lower and upper bounds, but also for all v between the VS bid–ask bounds and all w between the WTA bid–ask bounds. Figure 3 shows how our methods handle such situations. The distributions are asymmetric, as must be the case. Our methods accommodate $v < 0.5 < w$ with a negatively skewed distribution and $w < 0.5 < v$ with a positively
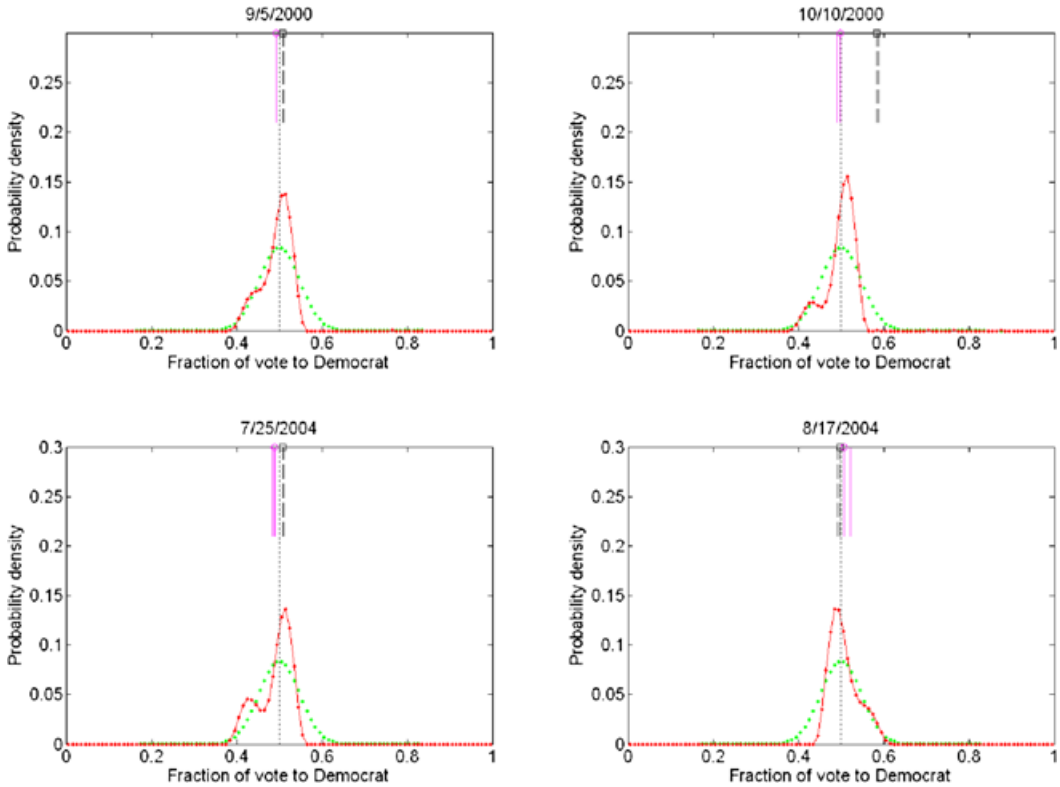
FIGURE 3. Estimated probability distribution on four sample days in which bound ranges for vote-share and winner-takes-all contracts were on opposite sides of 0.5.

skewed distribution. In all four cases, $\widehat{\mathbf{p}}_t$ brings both $E(f)$ and $P(f > 0.5)$ as close to 0.5 as possible, as indicated at the top of each panel.

Our procedure can also generate distributions with multiple modes. Two of the distributions in Figure 3 are unimodal, and two are bimodal. For some days, the distributions are multimodal. Figure 4 shows a multimodal distribution from 1992, a bimodal distribution from 2000, and the distribution on the following day in both cases.

Studying why multiple modes arise and disappear is informative. On September 22, 1992, the VS contract bounds were 0.535 and 0.540, and the WTA bounds were both 0.700. The indifferent trader's distribution (Figure 4, top left panel) places $E(f) = 0.535$, so the VS lower bound is binding. Hence, relative to the smooth concentrated distributions preferred by the prior, the VS range is too high. The indifferent trader's probability distribution reconciles this using the mass around the right mode. At 0.645, this mode is much higher than the range of any of the INT contracts (Table 1). The left mode reflects a lower bound for the INT contract with positive liquidation value if and only if $f < 0.46$. Given the high VS lower bound, the distribution assigns a mode just below 0.46, rather than a smoother distribution below 0.46 (which would decrease $E(f)$). The following day the VS contract range fell to [0.500, 0.533]. This drop completely resolves the conflict between the VS contract and the WTA contract whose September 23 range is [0.677, 0.720].
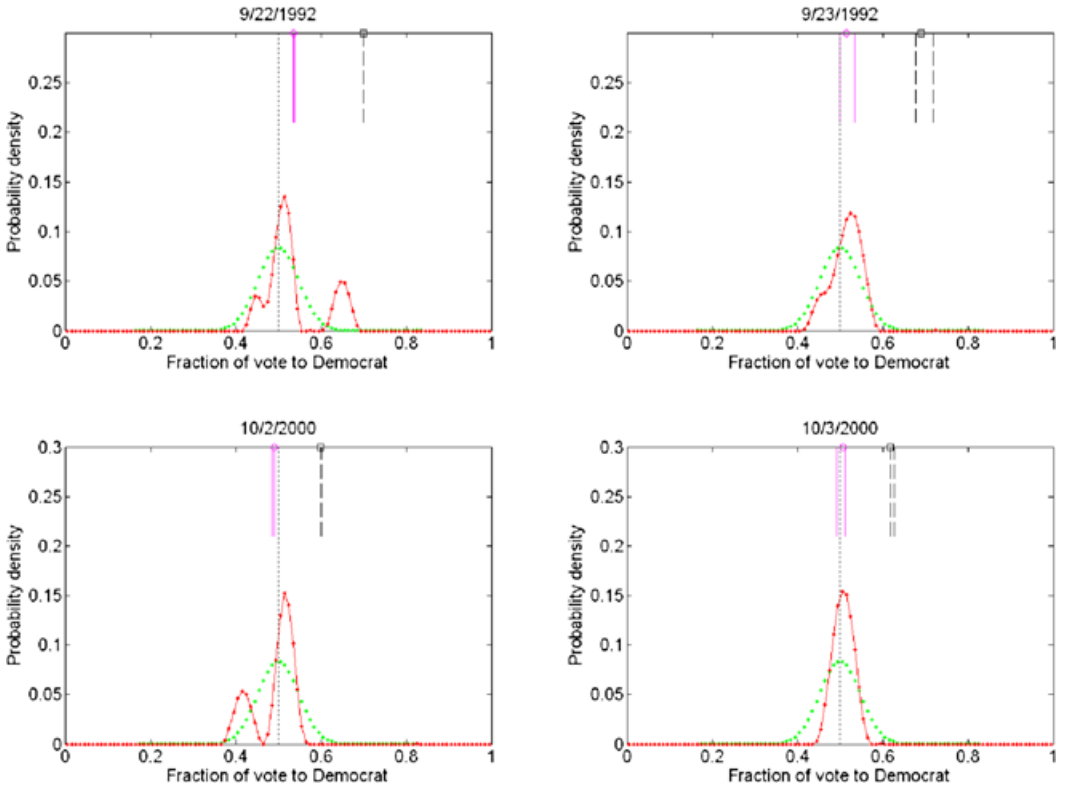
FIGURE 4. Estimated probability distribution on two sets of successive days in which there were large changes in the distributions.

None of the VS and WTA contract bounds remains binding and the right mode vanishes. While the left mode has also disappeared due to the drop in $E(f)$, the asymmetry attests to a relatively high lower bound for the INT contract paying off when $f < 0.46$.

Sometimes, multiple modes arise even without multiple interval contracts. In 2000, only VS and WTA contracts (and no INT contracts) traded. The bound configuration on October 2 (lower left panel, Figure 4) violates constraints from all of the specific distributions studied in Section 3. The left mode in that day's distribution results from a VS upper bound that is too low, relative to the WTA lower bound, in the context of the prior preferences for smooth concentrated distributions. Accordingly, the VS upper bound and the WTA lower bound are binding. The following day the same two bounds are binding. However, the VS upper bound has risen from 0.490 to 0.512. The WTA lower bound increased slightly more from 0.602 to 0.628. Yet this removed the bimodality in the distribution. The high concentration near the mean of the distribution is now consistent with the bounds of the WTA contract. In general, distributions are more sensitive to VS contract bounds than to WTA or INT bounds. Section 5.3 provides a more complete analysis of these issues by examining closely the role of upper and lower bounds in inferring the underlying distribution.

Changes between unimodality and bimodality occur regularly in all four years. The propensity for multimodality is generally stronger in the 1992 and 1996 elections than

in 2000 and 2004. In each election, there is substantial serial correlation in the number of modes exhibited in the distributions of the indifferent trader.

5.2 *The evolution of distributions*

The previous section gives two examples of how the indifferent trader's distributions evolve through time. More generally, the probability distributions $\widehat{\mathbf{p}}_t$ exhibit periods of several weeks with little change, punctuated by periods of similar length with much greater changes. The day-to-day change can be summarized by $\sum_{i=1}^{100} |\widehat{p}_{t,i} - \widehat{p}_{t-1,i}|$. This function is bounded below by 0 (corresponding to no change in the distribution) and above by 2 (corresponding to disjoint supports for the distribution on successive days). Figure 5 plots these changes as dots connected by a solid line for each of the 4 years. The horizontal line indicates the average change in each year: 0.216 in 1992, 0.092 in 1996, 0.117 in 2000, and 0.112 in 2004. (Vertical lines indicate dates when additional INT contracts were introduced; see Table 1.) In each case, the volatility of the distribution is persistent and there is no systematic tendency for the distribution to become more or less volatile as the election approaches. There is also no systematic relationship between the magnitude of changes in $\widehat{\mathbf{p}}_t$ and the sizes of bid–ask spreads in the VS and WTA markets.
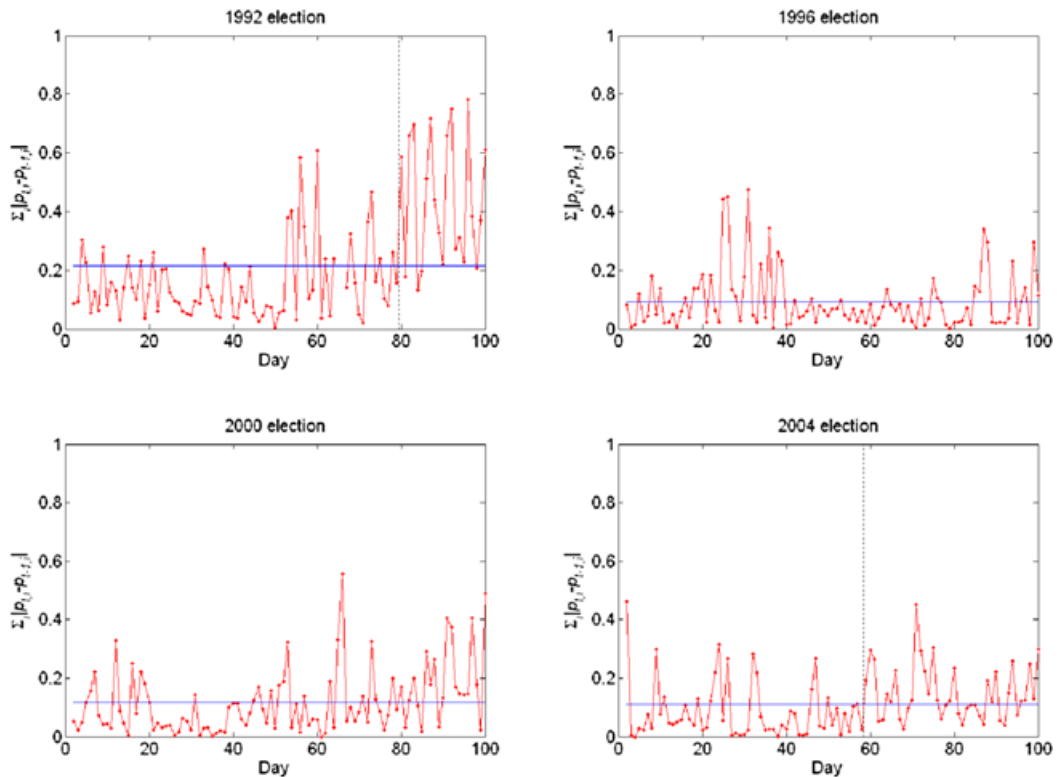

FIGURE 5. Daily changes in the estimated probability distributions.

5.3 *Analysis of constraints*

As detailed in the Appendix, combinations of contract bid and ask prices provide lower and upper bounds for a function of $\mathbf{p}_t$ specific to each contract. (Figures of the bounds appear in the Appendix.) Each day, our approach indicates whether the upper, lower, or neither bound is binding for $\widehat{\mathbf{p}}_t$. Typically there are many days in succession in which the upper, lower, or neither bound is binding. Clearly, the elections differ from each other. In 1996, the value of $P(f > 0.5)$ implicit in $\widehat{\mathbf{p}}_t$ is as high as it possibly can be (i.e., the upper bound of the WTA contract is binding) in all 100 days. For 97 of the 100 days in 1996, the lower bound of the VS contract is binding, making the value of $E(f)$ implicit in $\widehat{\mathbf{p}}_t$ as low as it can be on those days. On the other 3 days in 1996, neither VS bound was binding. In 2000 and 2004, upper bounds of VS contracts and lower bounds of WTA contracts were predominantly binding. In 1992, VS lower bounds were binding on 21 days and upper bounds were binding on 47 days, while WTA lower bounds were binding on 55 days and upper bounds were binding on 39 days.

If a lower (upper) bound is binding, then a hypothetical decrease (increase) in that bound would lead to $\widehat{\mathbf{p}}_t$ that has higher probability under our prior distribution—that is, the corresponding probability distribution would be smoother and/or more concentrated. A standard measure of the strength of a binding constraint is the derivative of the objective function (in our case, the log of the posterior density) with respect to a change in that constraint. For our quadratic objective function, this derivative is measured by the Lagrange multiplier ("shadow price") of the constraint in the solution of the quadratic programming problem, computed as a by-product of the solution. For the VS contracts, a Lagrange multiplier is positive if and only if a corresponding contract bound is binding. The same is true for WTA contracts on those days when no other INT contracts are traded. Since the bounds reflect sums of contract bounds on other days, Lagrange multipliers can be positive even if the individual contract bounds do not bind $P(f > 0.5)$; in these cases, the contract bounds are binding for some other function, for example, $P(f > 0.54)$ in 1992 or $P(f > 0.52)$ on the 42 days preceding the 2004 election.

Generally higher Lagrange multipliers for VS contracts arise because the distribution is much more sensitive to changes in $E(f)$ than it is to $P(f > 0.5)$ in the IEM presidential election markets. The sensitivity results from the concentration of the distribution—the fact that most of its support is typically over an interval of about 0.12 units. Figures 2–4 illustrate this fact and Section 5.4 provides more systematic evidence based on the standard deviation of the inferred distributions.[9] The large changes in $\widehat{\mathbf{p}}_t$ between successive days shown in Figure 4 are driven by large changes in binding VS contract bounds, while the WTA bounds changed relatively little. In general, the opposite must be true and, in general, it is, as indicated by the (v, w) means of lower and upper bounds shown in these figures.

---

[9]To understand the relative difference in sensitivity, consider a uniform distribution on an interval of length 0.125, and therefore height 8, that includes the point 0.5. If the mean of this distribution $E(f)$ increases by 0.01, then (as long as 0.5 remains in its support) $P(f > 0.5)$ increases by $0.01 \times 8 = 0.08$. Other distributions lead to variants on this result, but as long as the distribution is concentrated in a neighborhood of a value of f corresponding to an interval contract price, and movements in the distribution are primarily location shifts, then interval contract prices should move much more than VS contract prices.

   Preliminary analysis suggests that the Lagrange multipliers contain important information. Typically, the multipliers are positive on one side of the VS market and on the opposite side of the WTA market at any given time. For example, throughout the 1996 race, multipliers were typically positive on the lower bound of the VS market and the upper bound of the WTA market. This indicates that a lower VS bid and a higher WTA ask would be consistent with a smoother and more concentrated distribution. Indeed, there is a significant correlation between high VS Lagrange multipliers (on the lower bound) and decreases in VS prices over the next 24 hour period. While a complete analysis of their impact is beyond the scope of this paper, we point this out here to show that, through the Lagrange multipliers, our procedure generates additional, potentially useful, information.

5.4 Summary of the distributions

Figure 6 summarizes the distributions $\widehat{\mathbf{p}}_t$ of our indifferent trader over the 100 days preceding each of the four elections. The horizontal axis numbers the days before the elections, just as in Figure 5. The corresponding lines provide aspects of each day's distribution, the same distributions detailed in Figures 2, 3, and 4 for some specific days. These
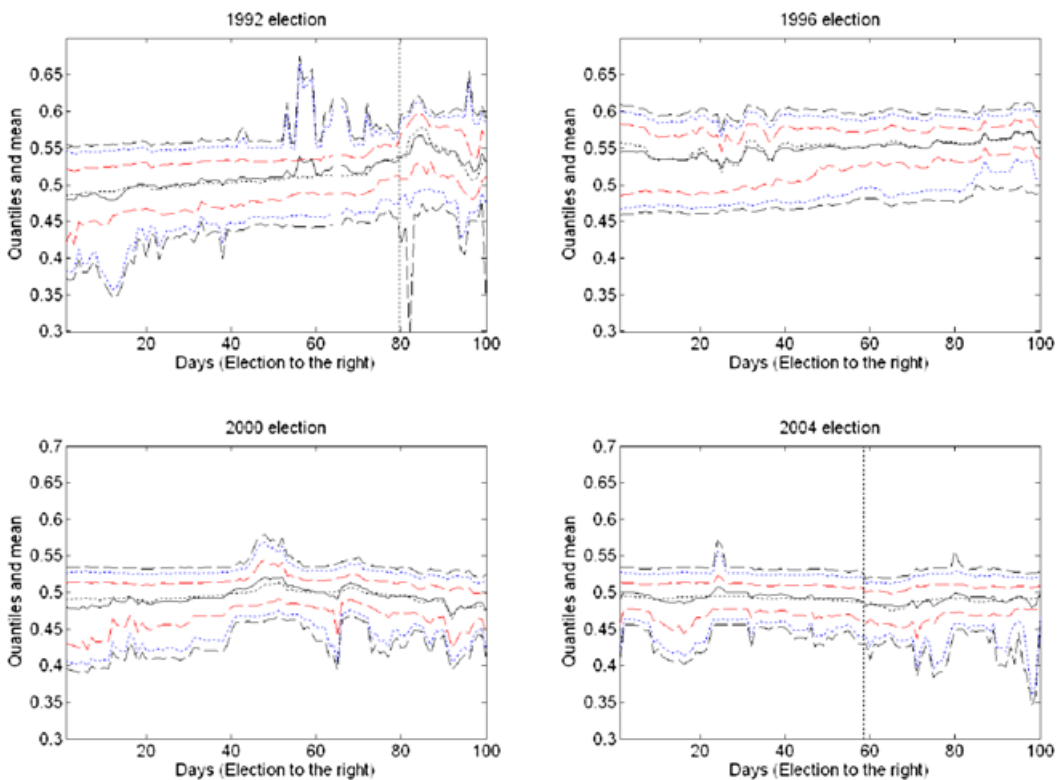

FIGURE 6. Solid lines are the means of the estimated distributions. Dotted lines near the means are the medians; light dashed lines denote the 0.25 and 0.75 quantiles; dotted lines denote the 0.10 and 0.90 quantiles; dark dashed lines denote the 0.05 and 0.95 quantiles.

aspects are indicated by a series of solid, dashed, and dotted lines. The solid line near the middle of the others is the mean of each day's distribution. The dotted line near the solid line is the median of each day's distribution. The dashed line immediately below is the first quartile and the one above is the third quartile. The dashed line at the top is the 0.99 quantile and the dotted line immediately below is the 0.95 quantile. The dashed line at the bottom is the 0.01 quantile and the dotted line immediately above is the 0.05 quantile.[10]

Periods of transient asymmetry are prominent features of the evolution of the distributions. Note the strong outward movement in the right tail of the distribution in days 56–59 (October 19–22) of the 1992 election, the virtual collapse of the right side of the distribution in days 20 and 21 (August 18 and 19) in 1996, and the tendency for a long left tail to emerge sporadically in 2004. In all of these cases, there is a corresponding shift in the mean, but the median is relatively unaffected.

The overall means and variances of the distributions suggest that the markets are informative. Figure 2 shows the distributions $\widehat{\mathbf{p}}_t$ the night before each election. The spreads in the distributions, difficult to see in the quantiles in Figure 6, are measured by their standard deviations in Figure 7. For reference, the panels in Figure 7 also show, as a light horizontal dashed line, the standard deviation of a logistic normal distribution fit to the Democratic share of the two-party popular vote in the 1868–2004 presidential elections. There is a lower bound to the standard deviation that corresponds to the unconstrained maximum of the prior density (i.e., the lower bound is achieved if no bounds are binding). As shown by the heavy dashed line, the lower bound is about half the standard deviation of the logistic normal distribution fit to the historical data. Frequently, the lower bound is approached in 2000 and 2004.[11]

Across all four elections, the standard deviation of our fitted subjective distribution is closer to this lower bound than it is to the historical standard deviation on more days than not.[12] However, there appears to be no systematic tendency for the spreads of the distributions to decrease as the election approaches.[13]

5.5 *Reliability of the distributions*

The changes in contracts traded in 1992 and 2004, at the points indicated by the vertical dotted lines in Figures 6 and 7, provide an informal indication of the reasonableness of our prior distribution. Since an increase in the number of contracts provides more

---

[10]In computing quantiles, the discrete distribution on the 100 points $0.005, \ldots, 0.995$ was spread evenly on the surrounding interval of length 0.01. This does not affect the mean of the distribution and, by making the support of the quantiles continuous rather than discrete, makes the chart easier to read.

[11]The lower bound was actually attained on one day in 2000 (day 71, 10/8/00) and three days in 2004 (day 34, 8/27; day 37, 8/30; and day 40, 9/1).

[12]There are exceptions, and in 1992 there are a few days when the standard deviation exceeds that of the reference historical distribution.

[13]The behavior of the distributions in 1996 after day 40 might bear that interpretation, but 1992 and 2004 might bear the opposite interpretation. None of these tendencies is very pronounced, and there is little indication that the subjective distributions of our hypothetical indifferent trader increases in precision as the election approaches.
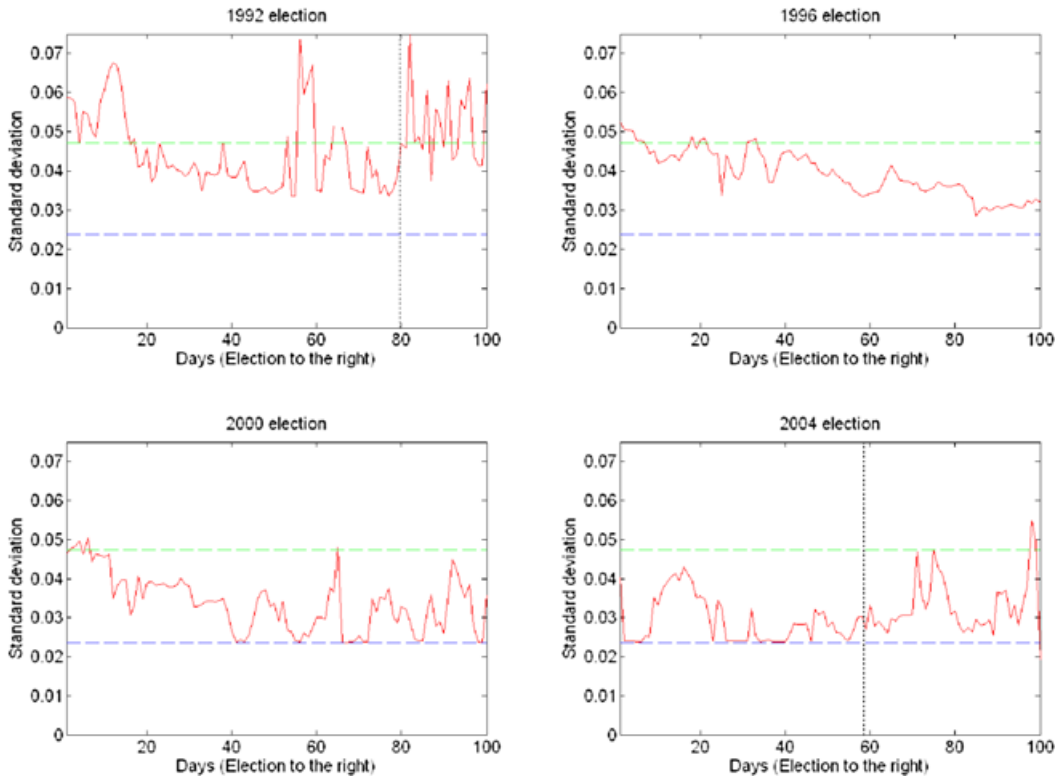
FIGURE 7. The solid line indicates the standard deviations of the estimated probability distributions. The light dashed line indicates the standard deviation of the historical f from 1868 to 2004. The heavy lower dashed line indicates the lowest possible standard deviation: that of the distribution corresponding to the unconstrained mode of the prior.

detailed information about the underlying distribution, these events have the potential to substantially affect inferences about the distribution. For example, the new contracts could indicate that the distribution had larger or smaller dispersion than previously believed, or that it was smoother or rougher. The results in Figures 6 and 7 hint at the possibility of such effects, but the changes near those dates are reflected in similar behavior at other dates in the same and other years. Our interpretation is that these events, which could raise serious questions about the model, fail to do so.

6. Conclusion

Prediction markets for future events are becoming increasingly common and are drawing sustained interest in the academic literature. Often there are several markets for the same event. In the case of the IEM presidential election markets, the focus of this paper, both vote-share and winner-takes-all contracts are traded, and on some days additional contracts are traded as well. While these markets were designed specifically for forecasting, given their proliferation, it is natural to ask whether the information in prediction

market prices can be integrated in a way that is useful more generally in decision making.

The approach taken in this paper is to derive a probability distribution for the event that is consistent with the prediction market prices. A probability distribution is a natural mode of integration of information, especially when there are multiple markets for the same event. In classical decision theory, uncertainty is represented by probability distributions, and so this mode of integration opens up prediction markets to analytical tools that are well understood in formal decision making.

In pursuing this approach, we followed two core principles. The first is that the probability distribution we seek is that of an indifferent trader: the subjective distribution of a normative risk-neutral individual who, given any portfolio of contracts traded in the prediction market, would choose not to reallocate that portfolio of contracts even if transactions costs were zero. In Proposition 1, we derive necessary and sufficient conditions for such a distribution to exist. These conditions are satisfied in the hundreds of configurations of IEM contract bid and ask prices that we study, and it is likely that these conditions are satisfied in other prediction markets in which arbitrage opportunities are quickly extinguished, as they are in the IEM. Such distributions are generally not unique, however. This characteristic is addressed by our second core principle, which is that the subjective distribution in question is the one that has the highest prior probability over all possible distributions consistent with market data, based on an explicit prior.

In general, the probability distribution of the indifferent trader need not belong to a conventional parametric family, like a beta or logistic normal distribution; neither need it have specific properties like symmetry or unimodality. Propositions 2–5 demonstrate that all of these conditions place constraints on contract prices that are violated in the IEM presidential election markets, and this outcome is likely in other prediction markets as well.[14] Thus our two principles imply abandonment of a parametric approach, and in this paper we have devised practical and reliable Bayesian nonparametric methods to infer the probability distribution of the indifferent trader.

The specifics of this approach were dictated by our desire to update the distribution of the indifferent trader in real time during the 2008 IEM presidential election markets. This led us to choose a functional form for the prior distribution such that determination of the posterior mode is equivalent to the solution of a conventional quadratic programming problem. With this solution, we updated the distribution of the indifferent trader every minute for several months leading up to the 2008 presidential election. Each computation required less than a second and there were no convergence failures of the quadratic programming algorithm in over 100,000 successive executions.

The application of this method to contract bid and ask prices leading up to the 1992–2004 presidential elections in this paper leads to subjective probability distributions for the indifferent trader that have several notable characteristics:

(i) They are substantially more concentrated than is the empirical distribution of popular vote shares of the major party candidates in the 1868–1988 presidential elec-

---

[14]For example, Figure 2 in Berg, Neumann, and Rietz (2009) often shows a narrow interior band between two wider ones, indicating the potential for a multimodal distribution.

tions. Relative to this distribution, contract prices in the IEM presidential election markets are informative.

(ii) The probability distributions display volatility similar to that found in financial asset prices, and with no systematic increase or decrease in volatility as election day approaches.

(iii) There is no systematic tendency for the indifferent trader's distribution to become more concentrated as election day approaches. This outcome is more consistent with Bayesian learning about a time-varying latent variable (preferences of the electorate, for voting or not voting as well as for candidates) than with prediction of a time series repeatedly observed (which, taken literally, would require that the election be held repeatedly every day).

(iv) The inferred distributions are bimodal on many days and multimodal on some. As detailed in Section 5, most of these occurrences can be interpreted in terms of the inconsistency of WTA and VS contract prices with a common smooth and concentrated distribution. Our hypothetical indifferent trader's distribution resolves WTA and VS prices in terms of asymmetric or multimodal distributions.

The technique developed in this paper applies to prediction markets with multiple contracts linked to a single underlying continuous distribution. There are many such markets run on the IEM and on other prediction markets around the world. For example, the IEM has organized the Google IPO markets (Berg, Neumann, and Rietz (2009)) and movie box office markets (Gruca, Berg, and Cipriano (2008)). Other examples include the Hollywood Stock Exchange's "MovieStocks,"[15] sales markets at Hewlett–Packard (Chen and Plott (2002)), and many Intrade financial markets, for example, the Dow Jones Index Markets.[16] Wolfers and Zitzewitz (2004), Tziralis and Tatsiopoulos (2007), and Berg, Forsythe, Nelson, and Rietz (2008) provide further examples.

References

Berg, J. E., R. Forsythe, F. D. Nelson, and T. A. Rietz (2008), "Results from a dozen years of election futures markets research." In *The Handbook of Experimental Economics Results* (C. R. Plott and V. L. Smith, eds.), 742–751, Elsevier Science, Amsterdam. [184]

Berg, J. E., R. Forsythe, and T. A. Rietz (1997), "The Iowa electronic market." In *Blackwell Encyclopedic Dictionary of Finance* (D. Paxson and D. Woods, eds.), Blackwell, Oxford. [166]

Berg, J. E., J. Geweke, and T. A. Rietz (2010), "Supplement to 'Memoirs of an indifferent trader: Estimating forecast distributions from prediction markets': Appendix." *Quantitative Economics Supplemental Material*, 1, http://qeconomics.org/supp/6/supplement.pdf. [166, 167]

Berg, J. E., F. D. Nelson, and T. A. Rietz (2008), "Prediction market accuracy in the long run." *International Journal of Forecasting*, 24, 283–298. [164, 165, 169]

---

[15] http://www.hsx.com/help/?page=glossary#moviestocks.
[16] http://www.intrade.com/.

Berg, J. E., G. R. Neumann, and T. A. Rietz (2009), "Searching for Google's value: Using prediction markets to forecast market capitalization prior to an initial public offering." *Management Science*, 55, 348–361. [164, 165, 183, 184]

Bernardo, J. M. and A. F. M. Smith (1993), *Bayesian Theory*. Wiley, Chichester, U.K. [174]

Chen, K. Y. and C. R. Plott (2002), "Information aggregation mechanisms: Concept, design and field implementation." Social Science Working Paper 1131, California Institute of Technology, Pasadena. [165, 184]

Forsythe, R., F. Nelson, G. R. Neumann, and J. Wright (1992), "Anatomy of an experimental political stock market." *American Economic Review*, 82, 1142–1161. [166]

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. [166]

Geweke, J. (1991), "Efficient simulation from the multivariate normal and Student-*t* distributions subject to linear constraints." In *Computing Science and Statistics: Proceedings of the Twenty-Third Symposium on the Interface* (E. M. Keramidas, ed.), 571–578, Interface Foundation of North America Inc., Fairfax. [174]

Geweke, J. (2005), *Contemporary Bayesian Econometrics and Statistics*. Wiley, Hoboken, New Jersey. [173, 174]

Gruca, T. S., J. Berg, and M. Cipriano (2008), "Incentive and accuracy issues in movie prediction markets." *The Journal of Prediction Markets*, 2, 29–43. [164, 184]

Leigh, A. and J. Wolfers (2006), "Competing approaches to forecasting elections: Economic models, opinion polling and prediction markets." Working Paper 12053, National Bureau of Economic Research. [164]

Manski, C. F. (2006), "Interpreting the predictions of prediction markets." *Economics Letters*, 91, 425–429. [164]

Oliven, K. and T. A. Rietz (2004), "Suckers are born, but markets are made: Individual rationality, arbitrage and market efficiency on an electronic futures market." *Management Science*, 50, 336–351. [166, 171]

Page, L. (2008), "Comparing prediction market prices and opinion polls in political elections." *The Journal of Prediction Markets*, 2, 91–97. [164]

Shiller, R. J. (1984), "Smoothness priors and nonlinear regression." *Journal of the American Statistical Association*, 79, 609–615. [172]

Tzirialis, G. and I. Tatsiopoulos (2007), "Prediction markets: An extended literature review." *The Journal of Prediction Markets*, 1, 75–91. [184]

Wolfers, J. and E. Zitzewitz (2004), "Prediction markets." *The Journal of Economic Perspectives*, 18, 107–126. [184]

Wolfers, J. and E. Zitzewitz (2006), "Interpreting prediction market prices as probabilities." IZA Discussion Paper 2092, Bonn. [164]

Zellner, A. (1971), *An Introduction to Bayesian Inference in Econometrics*. Wiley, New York. [169]