Price dynamics in political prediction markets

Saikat Ray Majumder[a,b], Daniel Diermeier[a,b], Thomas A. Rietz[c], and Luís A. Nunes Amaral[a,d,1]

[a]Northwestern Institute on Complex Systems, Northwestern University, Evanston, IL 60208; [b]Kellogg School of Management, Northwestern University, Evanston, IL 60208; [c]Henry B. Tippie College of Business, University of Iowa, Iowa City, IA 52242; and [d]Department Chemical and Biological Engineering, Northwestern University, Evanston, IL 60208

Prediction markets, in which contract prices are used to forecast future events, are increasingly applied to various domains ranging from political contests to scientific breakthroughs. However, the dynamics of such markets are not well understood. Here, we study the return dynamics of the oldest, most data-rich prediction markets, the Iowa Electronic Presidential Election "winner-takes-all" markets. As with other financial markets, we find uncorrelated returns, power-law decaying volatility correlations, and, usually, power-law decaying distributions of returns. However, unlike other financial markets, we find conditional diverging volatilities as the contract settlement date approaches. We propose a dynamic binary option model that captures all features of the empirical data and can potentially provide a tool with which one may extract true information events from a price time series.

econophysics | information aggregation | wisdom of the crowd | binary model

Prediction markets trade specifically designed futures contracts with payoffs tied to upcoming events of interest (1). A common type of prediction market contract is a binary option contract that pays off $1 if an event occurs and $0 otherwise. The contract design, which differentiates them from typical futures contracts, allows prices to be used as direct forecasts of event probabilities (2–5).

Although betting on election outcomes was quite common in the United States prior to the Second World War as discussed in ref. 6, emergence of modern prediction markets, with the goal of information aggregation and revelation, can largely be traced back to the markets introduced by the Iowa Electronic Markets (IEMs) in 1988 (7). Since then, prediction markets have been created for election outcomes (7), financial results of companies (8), scientific breakthroughs (9), incidence of infectious disease (10), geopolitical events (9), box office takes of movies (11), the outcomes of sporting events (12), and hurricane landfalls (13, 14), among others. They have also been proposed for topics ranging from terrorist attacks (15) to future Olympic sites (16). Hedge Street (17) now trades binary option contracts on gold, silver, crude oil, and foreign exchange. More significantly, the Chicago Board of Trade (CBOT) recently created binary options markets on the Federal Funds target rate (18), a leading indicator of the U.S. economy.

Given their accuracy, reaction speed, and data richness (3, 19–23), prediction markets provide researchers with the opportunity to precisely assess how external factors shape collective beliefs about the likelihood of a given event. Here, we consider the paradigmatic case of U.S. presidential elections. We use the tools of financial time series analysis and econophysics (24–26) to investigate the price dynamics of prediction markets with the goal of developing methods to identify the truly critical events during presidential campaigns. There are numerous known empirical regularities for price dynamics in stock, foreign exchange, commodity spot and futures markets (27–34). There is also some research on "ordinary" options returns (35, 36) and much on the relationship between options prices and stock returns (37–39). For details, refer to refs. 37 and 38, which survey the extensive literature on empirical option pricing research, stock options, options on stock indexes and stock index futures, and options on currencies and currency futures. However, empirical return characteristics for binary

options—which differ considerably from other financial instruments, including ordinary options contracts*—have not yet been documented.

As a first step toward our goal, we investigate the statistical properties of the prices in the two most active IEM presidential winner-takes-all markets. Our empirical analysis of the data for the Democratic contracts in year 2000 and Democratic and Republican contracts in 2004 reveals that the distribution of returns decays in the tail as a power law with an exponent $\alpha \approx 2.6$. However, for the Republican contracts in year 2000 we find that the return distribution decays as an exponential function with a characteristic decay scale $\beta \approx 0.9$. We conjecture that this may have resulted from the greater influence of partisan trading for this particular contract.

Our empirical analysis enables us to propose and test a dynamic binary options model with conditional jump sizes and diverging volatility. We demonstrate that the model reproduces all the main features of the price dynamics in binary option markets. The model also suggests a criterion for identifying extraordinary price movements arising in such markets due to significant information events and thereby raises the possibility that one may be able to identify those events that shape a political campaign.

Maturity of Prediction Markets

Prediction markets are a relatively new forecasting tool. Nonetheless, some markets have trade volumes similar to traditional futures markets. For example, the daily number of trades in the IEM electronic markets that we study is comparable to the number of trades for equity options for very large companies such as IBM or DELL on the New York Stock Exchange. In fact, the number of trades in the IEM Federal Funds market is much higher than that for the similar CBOT binary options on rate decisions by the U.S. Federal Reserve [see supporting information (SI) *Appendix* for details]. Thus, although the dollar value of the contracts traded in the IEM is small, they are very active markets. Moreover, experimental economics evidence (42) and evidence from the prediction markets themselves (21) show that, even for small monetary payoffs, active markets reveal trader information. These facts suggest that at least large prediction markets, such as IEM markets for U.S. presidential elections, are mature enough to warrant analysis.

Prediction markets have been remarkably successful in correctly predicting future events (3, 19, 21, 22). For example, in presidential elections prediction markets routinely outperform opinion polls (21). This generalizes to other domains as well (3, 22, 23). Moreover, prediction markets rapidly incorporate new information

---

*Binary options differ from ordinary options in three respects: (*i*) the payoff structure, (*ii*) the fact that there is no underlying traded asset, and (*iii*) pricing discontinuities at settlement, as we will show below.

as was demonstrated in the IEM "1996 Colin Powell Nomination market" (20) (see *SI Appendix* for details). Given their large trading volume, reaction speed, and accuracy, IEM, therefore, provides us with the opportunity to assess how external events shape collective beliefs about the likelihood of a given event in the context of a political campaign.

The Data

The IEMs are real money markets open 24 hours a day, 7 days a week with trading through the Internet. Trading on their own accounts, traders place "bids" to buy and "asks" to sell contracts. These orders are placed into price- and time-ordered queues. Traders may also set the expiration of the order. If no expiration is provided, the order is removed at 11:59 PM Central Standard Time (CST) the day after the order was placed in the queue. The highest bid and lowest ask are available to all traders logged into the market. Besides placing an order into the queue, a trader can also accept the best bid (ask) to buy (sell) a contract. All feasible trades are executed immediately.

The IEM records information on every trade, including whether the trade was executed at the bid or ask and whether there were multiple individual trades associated with a single order. For convenience, we build equal time-interval time series for price, number of trades, and volume in dollars, where the time interval is $\tau = 60$ sec. We have checked that the dynamics of equal-interval time series is similar to the time series with actual trade times.

The 2000 presidential election winner-takes-all market opened on May 1, 2000 with contracts associated with the Democratic, Reform, and Republican parties; the 2004 presidential election winner-takes-all market opened on June 1, 2004, with contracts associated with the Democratic and Republican parties. These markets traded binary options contracts tied to the election outcome (43, 44). Each traded contract was associated with a party and paid $1 if that party received the majority of the two-party or three-party popular vote.

In theory, traders in prediction markets price contracts according to their expectations, so the prices will be a noisy proxy for the aggregate estimated probability of the associated event[†]; see ref. 45 for a more detailed discussion. Thus, the price of the contract associated with the Democratic party indicates the probability (with some uncertainty) that the party's nominee will take the majority of the two-party vote. Note, however, that there will always be some residual uncertainty and, hence, prices should remain bounded away from $0 or $1 until settlement. For instance, in the 1996 IEM presidential winner-takes-all markets, months in advance of the election, it was forecast that Clinton would emerge as the winner. This was reflected in the prices of the Clinton contracts, which slowly approached, but never reached, $1.

Statistical Properties of the Returns

In the IEM presidential election markets, contracts are effectively settled on election day, which is well-known in advance: November 7 for 2000 and November 2 for 2004. We set the origin of the time axis at these settlement dates. The times in our time series are then indexed as
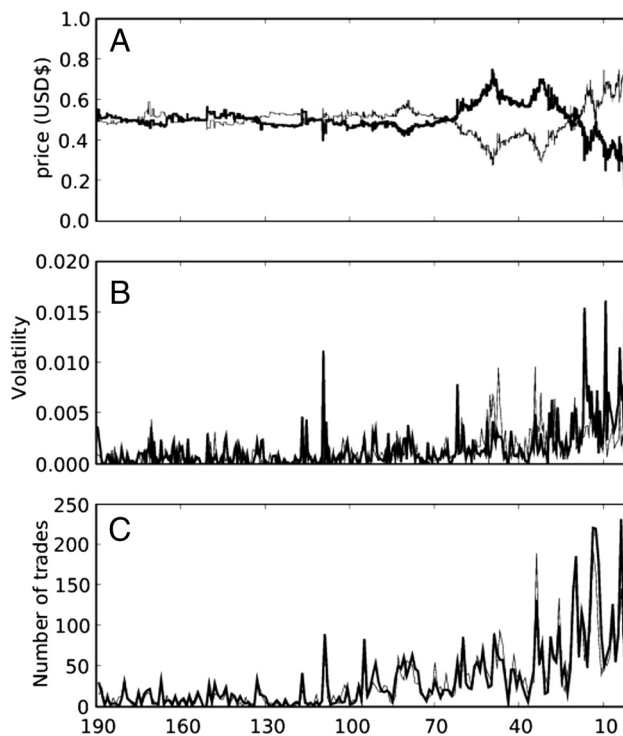
$$\tau_i = i\tau, \quad [1]$$

where $i = 0, 1, \ldots, N$, and $\tau_0 \equiv 0$ is the time when the contracts are settled and τ_N is when the market opens.

We define the return at time τ_i as

$$r(\tau_i) = \ln \frac{p(\tau_{i-1})}{p(\tau_i)}, \quad [2]$$

where $p(\tau_i)$ is the price of the contract at time τ_i.

―――――――――――――――

[†] Small amounts of noise may arise from the bid–ask spread, asynchronous trading, and "stale" prices, though such factors should be small in active markets.

Fig. 1. Trading dynamics in the 2000 presidential election market. (*A–C*) The bold line indicates the 2000 Democrat contract whereas the thin line indicates the 2000 Republican contract. Time is counted from the settlement date. (*A*) Price of the contracts in USD$. (*B*) Volatilities, estimated as standard deviation of the returns calculated on nonoverlapping 12-hour windows. (*C*) Daily number of trades. Note how the standard deviation of the returns and the number of trades both increase markedly toward the settlement date.

Although little is know about the price dynamics in prediction markets, there are three well-established facts about price fluctuations in stock markets, foreign exchange markets, and commodity markets (27–34). First, returns are uncorrelated for time scales longer than a few seconds. Second, volatilities are positively correlated over long time periods. Specifically, the correlations of the volatility decay as power laws. Third, the distribution of returns is consistent with a power-law asymptotic behavior,

$$P(|r| > x) \sim \frac{1}{x^\alpha}. \quad [3]$$

For stocks, foreign exchange rates, and commodity futures, the exponent $\alpha \sim 3$ (well outside the stable Lèvy regime $0 < \alpha < 2$) (24, 34), but $\alpha \sim 2.3$ for commodity spot prices (32).

We quantify the price dynamics of the Democratic and Republican contracts for the 2000 and 2004 elections along these three dimensions. We find that the number of trades increases dramatically toward the settlement date and that the returns in the final days of the market have significantly higher volatilities; cf. Fig. 1*C*. Specifically, conditional on a given price, the volatility is higher the closer the contract is to liquidation, that is, for a given $p(\tau_i)$, the volatility diverges as τ_i approaches zero (see *SI Appendix* for details). For this reason, we separately analyze the data in year 2000 for the final 10 days of the market (days 1-10), and for each of the previous two-month periods (days 11–70, 71–130, and 131–190). To avoid issues that may arise as information comes in on election day, we only analyze data up to midnight the day before the election [as is commonly done in the prediction market literature (3)].

To determine whether long-range correlations exist in the returns, we use detrended fluctuation analysis (46–48, 55), which works as follows. Consider a time series $x(t_i)$. One integrates this time series, generating a new time series $y(t_i)$, which is then divided

into blocks of size n. In each box, one performs a least-squares linear fit to the data (to capture any local trends at scale n), and determines the sum $F(n)$ of the squares of the residuals inside all the blocks of size n. This procedure is then repeated for different values of n. If $x(t_i)$ can be modeled as independent and identically distributed (i.i.d.) Gaussian variables, one finds

$$F(n) \sim n^{1/2}. \qquad [4]$$

Exponent values $>1/2$ indicate positive long-range correlations, whereas smaller values indicate long-range anticorrelations. For returns, we find an exponent ~ 0.5. For the volatilities, which we define here as the absolute value of the returns, we find an exponent ~ 0.7 (Fig. 2 A and B), except during the first two months of the market (days 131–190), when trading was very thin and the exponent is ~ 0.5. These results are consistent with the hypothesis that the returns display no correlations while there are positive long-range volatility correlations, similar to what is found in other financial markets.[‡]

Next, we estimate the power-law exponent α, defined in Eq. **3**, for the return distributions. As shown in Fig. 2C, the return distributions in days 1–10 are wider than for the previous months. However, we find that if we normalize the returns with the volatilities estimated separately in each one of the time periods, then the normalized return distributions follow the same functional forms. Specifically, the Kolmogorov–Smirnov (KS) test fails to reject the null hypothesis that the normalized returns are drawn from the same distribution.[§]

We compute the volatility for each one of the time periods as the standard deviation of returns over that time period,

$$\sigma_T^2 = \langle r(\tau_i)^2 \rangle_T - \langle r(\tau_i) \rangle_T^2, \qquad [5]$$

where T denotes one of the time periods and $\langle \ldots \rangle$ denotes a time average over the time period T. The normalized returns $\hat{r}_T(\tau_i)$ in T are then defined as

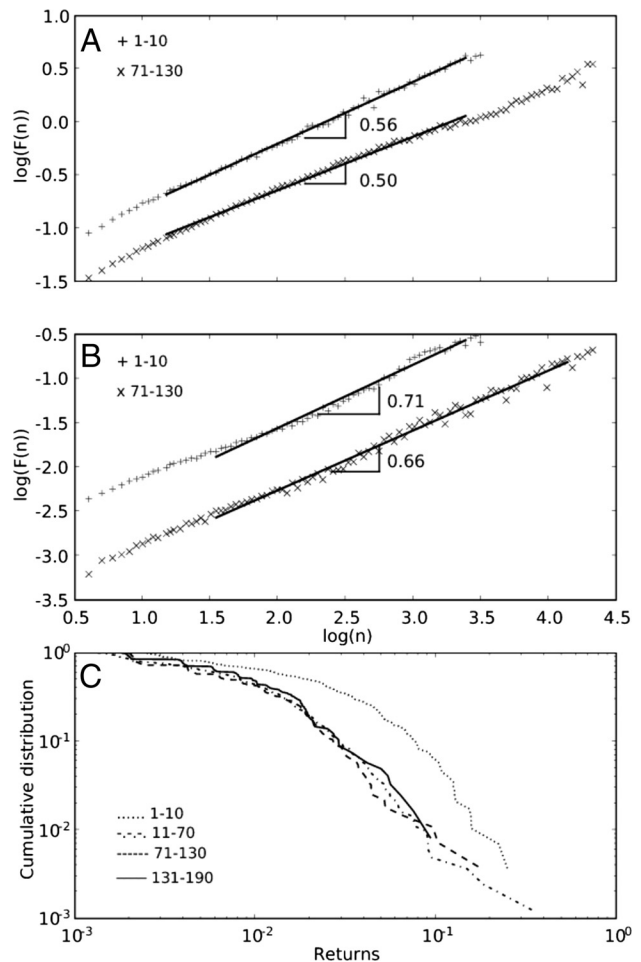$$\hat{r}_T(\tau_i) = \frac{r_T(\tau_i)}{\sigma_T}. \qquad [6]$$

Since these normalized return distributions have the same functional forms, we compute a single distribution for the positive and negative returns from the different time periods. Using the Hill estimator (54) and bootstrapping, we then get $\alpha = 2.6 \pm 0.2$ for the 2000 and 2004 Democratic contracts and the 2004 Republican contract (Fig. 3 A and C).[¶]

Surprisingly, for the 2000 Republican contract, we find the return distribution decays at an exponential rate,

$$P(|r| > x) \sim \exp(-\beta x), \qquad [7]$$

where β is the characteristic decay scale. We find that the tails of the return distributions decay with the rate $\beta = 0.9 \pm 0.1$ (Fig. 3 B and D). The fact that the Republican contracts are not perfectly negatively correlated with the Democratic contracts can be understood if one recalls that the market in 2000 included a Reform party winner-takes-all contract (in addition to the Democratic and Republican contracts).

The exponential decay of the return distribution can be attributed to partisan trading. For a well-functioning market in which traders have no partisan beliefs, one would expect traders to buy (sell) Democratic and Republican contracts at approximately

---

[‡] See *SI Appendix* for details of the method and the results. To make sure that those results are not due to the non-Gaussian distribution of the returns, we randomized the time order of the returns and reevaluated the exponent values. We find that, for the randomized time series, the exponent values ~ 0.5 for both the returns and volatilities.

[§] See *SI Appendix* for the P values from the KS tests. The confidence bounds in Fig. 3 A and B show that the deviations in the tails are consistent with expected fluctuations.

[¶] Refer to the *SI Appendix* for description of these and related statistical methods.

Majumder *et al.*

Fig. 2. Statistics of the returns for the 2000 Democratic contract. We use the detrended fluctuation method (see main text and *SI Appendix*) to quantify the correlations. (*A*) Autocorrelation of the returns. The data are consistent with uncorrelated returns. (*B*) Autocorrelation of the volatilities. We find power-law decaying positive correlations for the return volatilities. This implies that periods of large volatility are more likely to be followed by periods of large volatility than by periods of low volatility. (*C*) Distribution of the positive returns. The data suggest the possibility of an asymptotic power-law decay of the distribution of the returns.

equal rates. However, traders affiliated with a party tend to preferentially buy the contract of the party with which they are affiliated and to preferentially sell the contract of the other party. While the bias in those choices is relatively small for the 2000 and 2004 Democratic contracts and the 2004 Republican contract, they are stronger for the 2000 Republican contract. Relative to other contracts, more Republican traders in 2000 trade as if they truly believe the Republican candidate is going to win, and more Democrat traders trade as if they truly believe that the Republican candidate is going to lose. Thus, while those partisan Republican traders are very willing to buy the Republican contract, the partisan Democrat traders are very willing to sell it (see *SI Appendix*, Table 3).
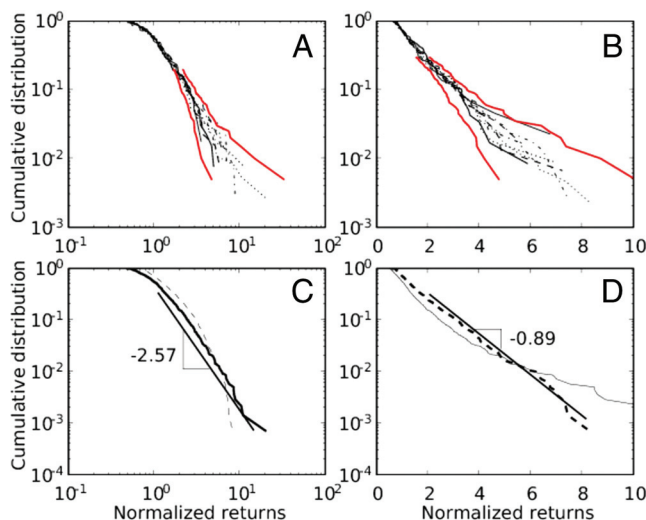
These biases have two consequences. First, these traders may take on substantial risk, since their portfolios will be heavily "tilted" toward one of the contracts. Second, partisan traders' inability to accommodate new information as rapidly as nonpartisan traders (49) results in their constant willingness to buy (or sell, depending on their bias) which prevents returns with larger magnitude from occurring. Interestingly, our findings for the 2000 Republican contract mirrors unexplained findings for the Indian stock market. Specifically, Matia *et al.* (50) reported an exponential decaying probability density function of the price fluctuations

Fig. 3. Asymptotic behavior of the distribution of returns for the 2000 presidential election market. (*A*) Double-logarithm plot of the distribution of normalized 2000 Democratic contract returns for the different time periods and for both the positive and the negative tails. The red lines show the 95% confidence intervals. This shows that the deviations in the tails of the distributions are consistent with the expected fluctuations. (*B*) Log-linear plot of the distribution of normalized 2000 Republican contract returns for the different time periods and for both the positive and the negative tails. The red lines show the 95% confidence intervals. This shows that the deviations in the tails of the distributions are consistent with the expected fluctuations. (*C*) Double-logarithm plot of the distributions of pooled normalized 2000 Democratic contract (full line) and of pooled normalized 2000 Republican contract (dashed line). A straight line in this plot indicates a power-law dependence. (*D*) Same data but in a log-linear plot. A straight line in this plot indicates an exponential dependence.

when they analyzed the daily returns for the period November 1994 to June 2002 for the 49 largest stocks of the National Stock Exchange in India. Our analysis suggests the hypothesis that a significant fraction of Indian traders may hold strong biases that determine their trading strategies.[∥]

The existing data do not allow a full exploration of this hypothesis at this stage. Thus, we can not give a full explanation of why the partisan effect was stronger in 2000 than in 2004. However, studies related to cognitive dissonance (52) or confirmation bias (53) would suggest that the effect would be stronger in elections with stronger emotional attachment to the respective candidates. This strikes us as a promising avenue for further research.

Model

Binary options liquidate at either $0 or $1. This implies a pricing discontinuity at maturity. The value of the option will jump from the current price to either $0 or $1 the instant the uncertainty is resolved. Another significant feature of binary option contracts is that the range of possible returns depends on the current price. For prices close to, for example, $1, the price can increase only by a very small amount, however, it can decrease by 100%. As a result, a plausible model must incorporate conditional asymmetric up and down jumps with increasing volatility as one approaches the settlement date.

Let T_a be the average time between consecutive trades and t_i the time at which the ith-to-last trade occurs. The median time difference between consecutive trades for the 2000 Democratic contract was ~60 sec and, therefore, we set $T_a = 60$ sec. We hypothesize that the current value of a winner-takes-all contract, which will settle at a value of $1 or $0, evolves according to

---

[∥] However, in ref. 56, R. K. Pan and S. Sinha have analyzed high-frequency tick-by-tick data for the Indian stock market and found that the cumulative distribution has a tail described by the power law with an exponent ~3 contrary to the findings in ref. 50.

$$\mathcal{P}(t_{i-1}) = \begin{cases} \mathcal{P}(t_i) + \frac{1 - \mathcal{P}(t_i)}{(t_i/T_a)^\gamma}, \text{with probability } \mathcal{P}(t_i) \\ \mathcal{P}(t_i) - \frac{\mathcal{P}(t_i)}{(t_i/T_a)^\gamma}, \text{ with probability } 1 - \mathcal{P}(t_i) \end{cases}, \quad [8]$$

where $\gamma \geq 0$ and $\mathcal{P}(t_i)$ refers to the value at the ith-to-last trade. This process is a Martingale; at any point in time $\langle \mathcal{P}(t_{i-1}) \rangle = \mathcal{P}(t_i)$.[**]

Additionally, we see that it converges to the appropriate value at settlement

$$\mathcal{P}_{\text{settlement}} = \begin{cases} 1, & \text{with probability } \mathcal{P}(t_1) \\ 0, & \text{with probability } 1 - \mathcal{P}(t_1) \end{cases} . \quad [9]$$

The model also makes it clear what we mean by conditional diverging volatility. The variance of returns from the underlying process is given by

$$\left\langle \left(\frac{\mathcal{P}(t_{i-1}) - \mathcal{P}(t_i)}{\mathcal{P}(t_i)} \right)^2 \right\rangle = \left(\frac{T_a}{t_i} \right)^{2\gamma} \left(\frac{1 - \mathcal{P}(t_i)}{\mathcal{P}(t_i)} \right). \quad [10]$$

This explicitly shows that, conditional on a given price, $\mathcal{P}(t_i)$, the volatility is expected to be higher the closer a contract is to liquidation.

Eq. **8** models the dynamics of the "true" value of the contract. The actual price, $p(t_i)$ will, however, deviate from $\mathcal{P}(t_i)$ due to noise and price information delays. To incorporate noise we model the observed price process in the following way

$$p(t_{i-1}) = \begin{cases} p(t_i) + \frac{1 - p(t_i) + \eta(t_i)}{(t_i/T_a)^\gamma} \exp[\epsilon(t_i)], \\ \quad \text{with probability } p(t_i), \\ p(t_i) - \frac{p(t_i) + \eta(t_i)}{(t_i/T_a)^\gamma} \exp[\epsilon(t_i)], \\ \quad \text{with probability } 1 - p(t_i) \end{cases} \quad [11]$$

where η and ϵ are Gaussian random variables. The additive noise term, η, prevents the prices $0 and $1 from becoming absorbing states of the dynamics. For η we model a Gaussian distributed variable with zero mean and a very small standard deviation; the results shown were obtained for a standard deviation of 0.0003. Because $\eta \neq 0$, the price, $p(t_i)$, deviates slightly from a martingale process.

By price information delays we refer to the fact that traders may not have access to the most current price but to a price some time in the past. Since there is a 15- to 30-sec time lag for the IEM trading system to update information we set the time lag in our model to 25 sec (20).

Another issue that also needs to be taken into consideration is that large-volume bids or asks that cross the opposing queue may not trade at a single price. Instead, they will "run" through the opposing queue generating a series of prices that all move in the same direction. We treat each such event as a single trade.

The model as defined above then has the following free parameters: the exponent γ and the mean (μ) and the standard deviation (σ) of the noise term ϵ. To estimate these parameters, we see from Eq. **11** that if we set $\delta p(t_i) = (p(t_{i-1}) - p(t_i))/(1 - p(t_i))$ for positive price changes and $\delta p(t_i) = |(p(t_{i-1}) - p(t_i))/p(t_i)|$ for negative changes (where $| \ldots |$ denote the absolute value) then we obtain,
$$[12]$$
$$\log \delta p(t_i) = \epsilon(t_i) - \gamma \log(t_i/T_a).$$

[**] Technically, $\mathcal{P}(t_i)$ is the risk neutral measure, but should approximate the true probability in the absence of significant hedging demand. This implies the best forecast of the next price is the current price. In fact, the current price is the best forecast of the settlement value. In context, this is Fama's weak form efficiency with a zero expected return (27). A continuous arbitrage opportunity built into the IEM restricts the risk-free rate to zero. The "unit portfolio" of both (or all three) contracts is risk free and can always be traded for $1 cash and vice versa. Cash accounts earn zero interest. Since the aggregate portfolio is also risk free, it earns a zero return and, hence, the returns to aggregate risk factors are zero. That is, all assets should earn the risk-free rate, in this case, 0. Pricing contingent claims with zero aggregate risk at expected value results from a simple extension of refs. 40 and 41.
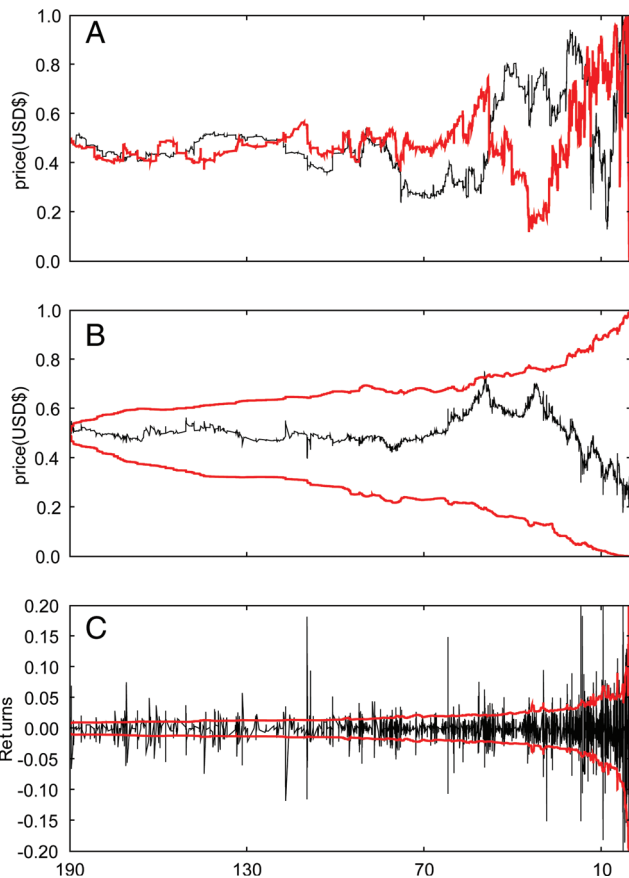
Fig. 4. The dynamic binary option model. (*A*) Two realizations of the binary option model. (*B*) Price of the 2000 Democratic contract (black line) and 90% confidence intervals for the model. (*C*) Returns of the 2000 Democratic contract (black line) and 90% confidence intervals for the model. Note how some returns in the data far exceed the 90% confidence intervals from the model, as may be expected given the power-law distribution of the returns.

We can estimate γ from the slope of the linear fit in Eq. **12** which can then be used to calculate σ from the standard deviation of the residuals.[††] We estimated $\gamma = 0.49 \pm 0.01$ and $\sigma = 1.22 \pm 0.02$ (refer to the *SI Appendix* for details).

We perform Monte Carlo simulations of the model with the estimated parameter values and find that the model generates uncorrelated returns and power-law decaying volatility correlations, in quantitative agreement with the empirical results. We also find that the actual price dynamics is well bounded by 90%

confidence bounds as shown in Fig. 4 *C* and *D* (for a description of this method and results from the model see *SI Appendix*). Additionally, we find that the distribution of returns decays as a power law. Using Hill estimator and bootstrapping, we estimate $\alpha = 2.3 \pm 0.2$, consistent with the estimate for the empirical data.

Discussion

The remarkable agreement between model predictions and the data may suggest a reasonably good understanding of the dynamics of prediction markets. However, there is one fundamental feature of prediction markets neglected by the model. In real prediction markets there is true information in the form of "known unknowns," such as the outcomes of debates or "unknown unknowns," such as revelations about the candidate's past, arriving at the market. These real information events can be viewed as exogenous processes and may be characterized by larger jumps than those arising from endogenous processes. It is then plausible that the identification of sharp differences between model predictions, the endogenous events, and real data, the exogenous events, could be used as a tool to identify information arrival at the market. In the context of a political contest, this approach can be used to determine which campaign events have a substantial impact on the fortunes of a particular candidate.

There is another possible application of our model which we believe will have a great impact in the course of a political campaign. In an election, there is a predetermined date when all the uncertainties are resolved, the settlement date. One may, however, realize that, in a particular election year, much of the uncertainty can be resolved earlier than the actual settlement date. For example, in the 1996 presidential election, it was forecast that Clinton would emerge as the winner about 100 days prior to the actual settlement date. Our model can be used to estimate this date by which most uncertainties are settled and, as a result, enable the political campaigners to judiciously assign their campaign resources.

Although our focus here is on political markets, our insights apply to binary options markets in general and thus will be important for traders, exchanges, regulators, policy makers, and forecasters alike. For example, our model can be used to forecast a distribution of likely price movements and, as a result, be used by exchanges to set margin requirements for traders of binary options conditional on prices and time to settlement. Another interesting aspect of our study is the possible application to crashes in financial markets. The approach to settlement date is remarkably similar to the increased volatility close to a market crash. Potentially, a generalization of our model could be used to estimate the time of a crash in these markets.

[††]We have assumed that ϵ has mean zero. One might, instead, assume that the mean is not zero and attempt to estimate it as well. In Eq. **11**, however, both the μ and γ effectively scale the jump sizes relative to the remaining time controlling the speed of convergence. As a result, they prove very difficult to identify independently without inordinate amounts of data. Preliminary analysis indicates a correspondence between the μ and γ estimates, where the speed of convergence weighs relatively more heavily on one parameter or the other. Pairs of estimates appear to explain the data equally. Here, we choose to model mean zero noise and let γ reflect the speed of convergence. We leave further exploration of the μ–γ relationship to future research.

1. Arrow KJ, *et al.* (2008) The promise of prediction markets. *Science* 320:877–888.
2. Hayek FA (1945) The use of knowledge in society. *Am Econ Rev* 35:519–530.
3. Wolfers J, Zitzewitz E (2004) Prediction markets. *J Econ Perspect* 18:107–126.
4. Wolfers J, Zitzewitz E (2006) Interpreting prediction market prices as probabilities. *Information Markets: A New Way of Making Decisions* (AEI Press, Washington, DC).
5. Tziralis G, Tatsiopoulos I (2007) Prediction markets: An extended literature review. *J Prediction Markets* 1:75–91.
6. Rhode P, Strumpf K (2004) Historical presidential betting markets. *J Econ Perspect* 18:127–141.
7. Forsythe R, Nelson FD, Neumann GR, Wright J (1992) Anatomy of an experimental political stock market. *Am Econ Rev* 82:1142–1161.
8. Bondarenko O, Bossaerts P (2000) Expectations and learning in Iowa. J Banking Finance 24:1535–1555.
9. Foresight exchange. Available at http://www.ideosphere.com. Accessed May 22, 2008.
10. Polgreen PM, Nelson FD, Neumann GR (2007) Use of prediction markets to forecast infectious disease activity. *Clin Infect Dis* 44:272–279.
11. Gruca T, Berg J, Cipriano M (2003) The Effect of electronic markets on forecasts of new product success. *Inf Syst Frontiers* 5:95.
12. Tradesports markets. Available at http://www.tradesports.com. Accessed May 22, 2008.
13. Huntley H (August 16, 2005) Put your money on the storm. *St Petersburg Times.* Available at http://www.sptimes.com/2005/08/16/Business/Put_your_money_on_the.shtml. Accessed December 27, 2008.
14. Chicago mercantile exchange hurricane futures. Available at http://www.cme.com/trading/prd/weather/hurricane.html. Accessed May 22, 2008.
15. Looney R (2003) Financing wars on terrorism and Iraq. *Strategic Insight* 2. Available at http://www.ccc.nps.navy.mil/si/jan03/terrorism.asp. Accessed December 27, 2008.
16. Barnhart B (April 17, 2007) Olympics venue an intriguing bet, just not in Chicago. *Chicago Tribune*, p 5.

APPLIED MATHEMATICS

POLITICAL SCIENCES

17. Hedge street. Available at http://www.hedgestreet.com. Accessed May 22, 2008.
18. Chicago board of trade, federal funds target rate. Available at http://www.cbot.com/. Accessed May 22, 2008.
19. Berg JE, Forsythe R, Nelson FD, Rietz TA (2008) Results from a dozen years of election futures markets research. *The Handbook of Experimental Economics Results* (Elsevier Science, Amsterdam).
20. Berg JE, Rietz TA (2006) The Iowa electronic markets: stylized facts and open issues. *Information Markets: A New Way of Making Decisions* (AEI Press, Washington, DC).
21. Berg J, Nelson F, Rietz T (2008) Prediction market accuracy in the long run. *Int J Forecast* 24:283–298.
22. Hahn RW, Tetlock PC (2006) *Information Markets: A New Way of Making Decisions* (AEI Press, Washington, DC).
23. Pennock DM, Lawrence S, Giles CL, Nielson F (2001) The real power of artificial markets. *Science* 291:987–988.
24. Mantegna RN, Stanley HE (2000) *Introduction to Econophysics: Correlations and Complexity in Finance* (Cambridge Univ Press, New York).
25. Farmer JD, Smith DE, Shubik M (2005) Is economics the next physical science? *Phys Today* 58:37–42.
26. Bouchaud J-P, Potters M (2003) *Theory of Financial Risk and Derivative Pricing: From Statistical Physics to Risk Management* (Cambridge Univ Press, New York).
27. Fama E (1970) Efficient capital markets: A review of theory and empirical work. *J Finance* 25:383–417.
28. Fama E (1991) Efficient capital markets: II. *J Finance* 46:1575–1618.
29. James D, Fama E, French K (2000) Characteristics, covariances, and average returns: 1929 to 1997. *J Finance* 55:389–406.
30. Gabaix X, Gopikrishnan P, Plerou V, Stanley HE (2006) Institutional investors and stock market volatility. *Q J Econ* 121:461–504.
31. Gabaix X, Gopikrishnan P, Plerou V, Stanley HE (2003) A theory of power-law distributions in financial market fluctuations. *Nature* 423:267–270.
32. Matia K, Amaral LAN, Goodwin SP, Stanley HE (2002) Different scaling behaviors of commodity spot and future prices. *Phys Rev E* 66:045103.
33. Gopikrishnan P, Plerou V, Amaral LAN, Meyer M, Stanley HE (1999) Scaling of the distribution of fluctuations of financial market indices. *Phys Rev E* 60:5305–5316.
34. Plerou V, Gopikrishnan P, Amaral LAN, Meyer M, Stanley HE (1999) Scaling of the distribution of price fluctuations of individual companies. *Phys Rev E* 60: 6519–6529.
35. Sheikh AM, Ronn EI (1994) A characterization of the daily and intraday behavior of returns on options. *J Finance* 49:557–579.
36. Coval JD, Shumway T (2001) Expected option returns. *J Finance* 56:983–1009.
37. Bates DS (1996) Testing option pricing models. *Handbook of Statistics* (Elsevier, Amsterdam), Vol 14, Chap 20, p 567.
38. Bates DS (2003) Empirical option pricing: A retrospection. *J Econometrics* 116:387.
39. Bakshi G, Cao C, Chen Z (2000) Do call prices and the underlying stock always move in the same direction? *Rev Finan Studies* 13:549–584.
40. Caspi Y (1974) Optimum allocation of risk in a market with many traders. *Allocation under Uncertainty: Equilibrium and Optimality* (MacMillan Press, London), pp 89–97.
41. Malinvaud E (1974) The allocation of individual risks in large markets. *Allocation under Uncertainty: Equilibrium and Optimality* (MacMillan Press, London), pp 110–125.
42. Sunder S (1995) *Handbook of Experimental Economics* (Princeton Univ Press, Princeton, NJ).
43. Arrow K J (1964) *The role of securities in the optimal allocation of risk-bearing*. Review of Economic Studies, Vol 31, pp 91–96.
44. Debreu G (1959) *Theory of Value: An Axiomatic Analysis of Economic Equilibrium* (Wiley, New York).
45. Wolfers J, Zitzewitz E (2006) *Interpreting Prediction Market Prices as Probabilities*. Federal Reserve Bank of San Francisco Working Paper Series (Federal Reserve Bank of San Francisco, San Francisco, CA).
46. Goldberger AL, *et al.* (2002) Fractal dynamics in physiology: Alterations with disease and aging. *Proc Natl Acad Sci USA* 99:2466–2472.
47. Cizeau P, Liu Y, Meyer M, Peng C-K, Stanley HE (1997) Volatility distribution in the S&P500 Stock Index. *Physica A* 2445:441.
48. Liu Y, *et al.* (1999) The statistical properties of the volatility of price fluctuations. *Phys Rev E* 60:1390.
49. Granberg D, Brent E (1983) When prophecy bends: The preference-expectation link in US presidential elections, 1952-1980. *J Personality Social Psychol* 45:477–491.
50. Matia K, Pal M, Salunkay H, Stanley HE (2004) Scale-dependent price fluctuations for the Indian stock market. *Europhys Lett* 66:909–914.
51. Forsythe R, Rietz TA, Ross TW (1999) Wishes, expectations and actions: a survey on price formation in election stock markets. *J Econ Behav Organ* 39:83–110.
52. Rabin M (1994) Cognitive dissonance and social change. *J Econ Behav Organ* 23:177–194.
53. Rabin M (1998) Psychology and economics. *J Econ Lit* 36:11–46.
54. Hill BM (1975) A simple general approach to inference about the tail of a distribution. *Ann Stat* 3:1163–1174.
55. Ashkenazy Y, *et al.* (2001) Magnitude and sign correlations in heartbeat fluctuations. *Phys Rev Lett* 9:1900.
56. Pan RK, Sinha S (2007) Self-organization of price fluctuation distribution in evolving markets. *Europhys Lett* 77:58004.