Information Processing in Prediction Markets: An Empirical Investigation*

Urmee Khan[†] and Robert P. Lieli[‡]

November, 2010

Abstract

This paper analyzes the direction and timing of information flow between prices, polls, and media coverage of events traded on prediction markets. We examine the race between Barack Obama and Hillary Clinton in the 2008 Democratic primaries for presidential nomination and ask the following questions: (i) Do prediction market prices have information that is not reflected in contemporaneous polls and media stories? (ii) Conversely, do prices react to information that appears to be news for pollsters or is prominently featured by the media? We construct time series variables that reflect the 'pollster's surprise' in each primary election as well as indices that describe the extent of media coverage received by the candidates. We carry out Granger causality tests between the day-to-

*Many thanks to Péter Kondor and Max Stinchcombe for helpful conversations. All errors are our responsibility. Preliminary—please do not cite without permission.

[†]Department of Economics, University of Texas, Austin.

[‡]Department of Economics, Central European University, Budapest and the National Bank of Hungary. The views expressed in this paper do not necessarily reflect those of the Bank.

day percentage change in the price of the 'Obama wins nomination' security and these information variables. There seems to be two-way Granger causality between price changes and the surprise element in the primary results. There is also evidence of one-way Granger causality in the short run from price changes towards media indices. These results suggest that prediction market prices anticipate at least some of the discrepancy between the actual outcome and the latest round of polls before the election. Nevertheless, prices also seem to be *driven* partly by election results, suggesting that there is an element of the pollster's surprise that is genuine news for the market. Furthermore, it seems prices capture information earlier than its revelation in the news media, although the apparent strength of this effect depends on modeling decisions such as the inclusion autoregressive conditional heteroskedasticity terms in the price process.

"We must look at the price system as such a mechanism for communicating information if we want to understand its real function"

-F.A. Hayek, *The Use of Knowledge in Society*, AER, 1945

1 Introduction

In the last few decades, asset markets designed for and exclusively dedicated to gathering information about probable outcomes of future events have come to existence. These markets, popularly known as prediction or information markets, have attracted attention for making accurate predictions about election outcomes [Berg et al., 2003, 2008a,b], product sales [Chen and Plott, 2002], film box office and myriad other variables of interest. With some operational variations, the fundamental design of such markets is as follows. The future event of interest is formalized as a random variable whose outcome would depend on an unobserved underlying state about which information is assumed to be dispersed among potential participants in the market. A set of complete contingent assets are introduced, whose eventual payoffs are tied to the future realization of the outcome of interest. The participants are expected to trade in the contingent assets based on their private information about the underlying state. Subject to assumptions about arbitrage possibilities and risk attitudes of the traders, the resultant market prices could be considered as a probabilistic prediction of the outcome of interest, at least in theory.

The performance of such markets in predicting the respective events of interest have created an optimism about its potential as a forecasting technology (Arrow et al. [2008], Sunstein [2007], Berg and Rietz [2003]). There is a nascent industry around this technology, and a growing stream of literature trying to understand various aspects of the information aggregation process in prediction markets [Tziralis and Tatsiopoulos, 2007]. Such comprehensive understanding would be vital for future improvement in this technology. Nevertheless, there remains a number of open questions, both theoretical [Wolfers and Zitzewitz, 2004] and empirical. The set of open questions include, but are not limited to, the following: i) What kind of information is built into the price? Do prices contain information beyond that collected by traditional mechanisms such as polls and the news media? ii) How fast does new information get built into the price? iii) What properties (statistical moments) of the price process best exhibit the impact of new information? iv) How does the market perform in terms of forecasting relative to other mechanisms (like polls)?

In comparing polls and prediction markets, the existing empirical literature has focused mainly on the last question. In a series of papers, Berg et al. have shown, using Iowa Electronic Market (IEM) data, that prediction market prices have significantly lower forecast errors both in the short and the long run compared to contemporaneous polls [Berg et al., 2003, 2008a,b]. In response, Erikson and Wlezien [2008] point out that since polls and prediction markets ask inherently different questions, they are not directly comparable. The standard question asked in a pre-election poll is this: ”if the election were held today, who would you vote for?” Therefore, polls only capture *voters' preferences at the moment* and should not be interpreted as forecasts in themselves; rather, sophisticated forecasts should build on the information provided by the polls. Presumably, markets do just that, as traders take into account their private and public sources of information before placing their bid, providing, in effect, an answer to the question: “who do you think will win”? Nevertheless, other sophisticated forecasts may exist. Using data from the same time period as Berg et al. [2008b], Erikson and Wlezien [2008] construct projections of vote shares and win probabilities based on daily polls and show that those projections can outperform prediction markets.

In contrast, our goal is to investigate the information flow between prediction market prices and two other conventional aggregators of information, namely media and opinion polls. We are not directly concerned with forecasting performance; rather, we seek to establish stylized facts that go towards answering questions i), ii) and iii). To this end, we examine the sequence of primary elections in the 2008 race for the Democratic nomination for president, and the evolution of prices in the market for the “Obama wins nomination” security. We construct time series variables that reflect the “pollster's surprise” in each primary election, measured as the difference between Obama's actual vote share and vote share predicted by the latest poll(s) before the primary, as well as indices that describe the extent of daily media coverage devoted to him. We conduct Granger causality tests between the daily percentage change in the price of the “Obama wins nomination” security and these information variables. These tests provide answers to the following operationalized versions of question i): Do prediction market prices have information that is not reflected in contemporaneous polls and media stories? Conversely, do prices react to information that appears to be news for pollsters or is prominently featured by the media? Further, the time horizons over which Granger causality relationships can be established provide answers to question ii). Finally, regarding question iii), the current version of the paper contains results about linear projections only. While under some auxiliary assumptions these projections can be interpreted as models of the conditional mean, we have not yet conducted tests aimed at detecting Granger causality in higher moments.

Our main empirical finding regarding the relationship between price changes and the pollster's surprise is strong two-way Granger causality. Thus, on the one hand, part of the pollster's surprise is predictable by previous price movements, so market prices appear to contain information not contained in the polls. On the other hand, the pollster's surprise also contains information about future price movements even when the history of prices is taken into account. Therefore, there is an element of the pollster's surprise that is genuine news for the market as well.

Regarding the relationship between prediction markets and media covarage, there is some evidence of one-way Granger causality in the short run from price changes towards some of the media news indices. Media coverage, particularly in the election season, contains both useful information and hype. Therefore, even if prediction markets are actually able to pick out the truly informative content from the media and discard the rest, the quantitative level of media coverage need not show a particularly strong empirical relationship with prices. Our baseline tests still deliver evidence that prices capture information earlier than its revelation in the news media, although the apparent strength of is sensitive to modeling decisions such as the inclusion autoregressive conditional heteroskedasticity terms in the price process.

2 Data

For prediction market prices, we use data from the IEM winner-take-all market for predicting the winner of the Obama vs. Clinton nomination race in 2008 for the period January 1, 2008 to June 15, 2008.[1] The raw data is a daily time series of prices (recorded every day at midnight, the single price being the last transaction price before midnight).

Data on state-based polls on primary outcomes was compiled from the website Real Clear Politics[2]. The polls are taken on the eve of the Democratic primaries in the respective states, conducted on likely Democratic voters, essentially asking the question who they would support in the upcoming primary (since polls were conducted by different organizations, the exact version of the question might have varied).

For each primary election, we consider the difference between Obama's actual vote share and his predicted vote share as measured by various polls in that state/territory. This raw difference is then multiplied by the fraction of total delegates at stake in the election (hence, a one percentage point difference in actual versus predicted vote share in a 'big' state counts as a larger surprise than the same observed difference in a 'small' state). To compute total surprise for a given day, we add the weighted surprise measures for all elections held that day. If there are no elections on the day in question, we set the surprise measure to zero.

[1] In the IEM winner-take-all market, one share of a candidate pays off 1 dollar if the candidate wins and nothing if the candidate loses. A portfolio of one unit of each candidate pays exactly 1 dollar. A trader who buys one unit of a candidate at, say 30 cents, wins either 1 dollar (a 70 cent profit) or nothing (a 30 cent loss) if the contract is held until market closing following the election. If the trader buys at 30 cents and sells at, say 70, the profit is 40 cents. For further details, consult the IEM website, http://www.biz.uiowa.edu/iem/.. .

[2] http://www.realclearpolitics.com

We measure the vote share predicted by polls in two different ways. First, we consider surprise relative to the average poll data published by the Real Clear Politics website. This average is taken over polls conducted by various polling agencies in the state in question over some period of time before the election. Second, we consider only the latest poll data available before the election (typically taken just one or two days before). If there is more than one such poll closing on the same day, then we take a weighted average with weights proportional to the sample size used in the polls. All raw poll data are taken from Real Clear Politics. There are a number of smaller states and territories where elections are held but for which no poll data is available. These are ignored altogether in constructing the surprise measures (i.e., their contribution to total surprise is set to zero).

For constructing the media variables, we use the data library of the Pew Research Centre Project for Excellence in Journalism (PEJ)[3]. They maintain a continually updated database of news stories in various kinds of media outlets which are monitored at a daily basis in regular intervals during the day. Stories are classified according to topic, lead newsmaker and placement prominence. Word counts for print and web stories and duration-in-seconds for broadcast/cable stories are also counted. For details about the sampling process, see Appendix.

For our study, we selected, from the PEJ daily sample for the relevant date range (January 1, 2008 to June 15, 2008), the stories that were categorized under the topic 'campaign', had Obama as the lead newsmaker, were news items as opposed to opinion pieces, and had relatively high prominence in their placements. For stories matching all the criteria above, we constructed time series variables representing total number of stories and total word counts per day for print media (Printcount and Wordcount respectively), total number of web

[3]http://www.journalism.org

stories (Webcount), total number of TV and radio stories (TV and Radio Count) and total duration-in-seconds for TV and Radio stories (Duration).

3 Methodology and baseline results

Let P_t denote the day t closing price of the Arrow-Debreau security that pays \$1 if Obama wins the Democratic nomination. As P_t exhibits a marked upward trend over the sample period, we work with the daily return series, denoted $p_t \equiv \log(P_t) - \log(P_{t-1})$, which does not show any apparent sign of being non-stationary. Further, let x_t denote any of the information variables (a measure of the pollster's surprise or a media index) described in Section 2.

A fairly narrow (but practical) definition of Granger causality states that x Granger causes p if the linear projection of p_t on $p_{t-1}, p_{t-2}, \dots$ differs from the linear projection of p_t on the larger information set $p_{t-1}, p_{t-2}, \dots, x_{t-1}, x_{t-2}, \dots$ (Granger [1969]). Assuming that the linear projections involved in the definition can be represented by a finite number of lags, one can readily test for causality going from x to p by estimating the model

$$p_t = \alpha_0 + \alpha_1 p_{t-1} + \ldots + \alpha_K p_{t-K} + \beta_1 x_{t-1} + \ldots + \beta_K x_{t-K} + \epsilon_t, \qquad (1)$$

and conducting a test of the hypothesis $H_0 : \beta_1 = \ldots = \beta_K = 0$.[4] In practice it is of course necessary to fix the lag length K. Given that observations are made at a daily frequency, and day of the week effects might be present, $K = 7$ does not, a priori, seem excessive. We in fact set the upper bound for K at 14 lags and report the results for $K = 1,2,3,7,10$. While small values of K (e.g., $K = 1,2,3$) are not sufficient to ensure that the residuals $\hat{\epsilon}_t$ are approximately white noise, these

[4] We opt for a quasi-likelihood ratio test (under the normality assumption), which is asymptotically equivalent to a Wald test.

tests still provide information about which lags of x have the most predictive power, i.e. over what time horizons is Granger causality present.

If one is willing to assume that, for some finite value of K, ϵ_t in equation (1) is a martingale difference sequence w.r.t. to its own history and the history of x, then a stronger interpretation of the corresponding test is available. In particular, the systematic part of equation (1) will represent the conditional mean of p_t given $p_{t-1}, p_{t-2}, \ldots, x_{t-1}, x_{t-2}, \ldots$, and the test becomes a test of Granger causality in the mean.

To test for causality in the reverse direction, i.e. going from p to x, one can consider interchanging the role of x and p in equation (1). However, the information variables constructed in Section 2 are nonnegative, and are given by a sequence of spikes with an interval of zero values in between. A finite order autoregressive model is unlikely to provide a good approximation to the projection of x_t on $x_{t-1}, x_{t-2}, \ldots, p_{t-1}, p_{t-2}, \ldots$. (Interpreting such a model as the conditional mean of x is especially problematic; see Engle [2002]) We circumvent this problem by implementing the Granger causality test as proposed by Geweke et al. [1983]. The method builds on the work of Sims [1972] and consists of regressing p_t on its own history as well as past, current, and *future* values of x:

$$p_t = \alpha_0 + \alpha_1 p_{t-1} + \ldots + \alpha_K p_{t-K} + \beta_1 x_{t-1} + \ldots + \beta_K x_{t-K} + \gamma_0 x_t + \gamma_1 x_{t+1} + \ldots \gamma_K x_{t+K} + \epsilon_t \tag{2}$$

Lack of Granger causality from p to x corresponds to the condition $H_0 : \gamma_1 = \ldots = \gamma_K = 0$. Using this equivalent formulation of Granger causality allows us to avoid building a model where x is a dependent variable rather than a regressor. Again, we consider models up to 14 lags/leads and report the results for $K = 1, 2, 3, 7, 10$.

The effective sample used for estimating all models described above ranges from January 15, 2008 to June 1, 2008 and consists of 139 observations. In addi-

Table 1: Does $\{p_t\}$ Granger cause $\{x_t\}$? Tests of H_0 : NO vs. H_1 : YES.

	Lags (K)				
Variable	1	2	3	7	10
Election surprise (RCP avg.)	0.006	0.002	0.002	0.000	0.000
Election surprise (latest poll)	0.037	0.034	0.037	0.000	0.000
Wordcount	0.019	0.047	0.080	0.047	0.338
Printcount	0.026	0.049	0.112	0.605	0.730
TV and radio count	0.697	0.700	0.689	0.667	0.679
Webcount	0.068	0.132	0.097	0.092	0.236
Duration	0.301	0.288	0.375	0.677	0.691

Note: The reported figures are p-values.

tion, fourteen pre-sample and fourteen post-sample observations are used to create lags and leads of p_t and x_t. Results for these benchmark Granger causality tests are displayed in Table 1 (from p to x) and in Table 2 (from x to p).

Looking at Table 1, there is strong Granger causality going from p_t to both measures of the pollster's surprise. That is, past price changes seem to be useful for predicting future election surprises (relative to polls) even when the history of surprises is taken into account. The result is very robust to the number of lags/leads included in model (2). Granger causality in this direction suggests that part of what appears to be a surprise for the pollster is already known to the market. On the other hand, in Table 2 we also find strong evidence of Granger causality from the pollster's surprise to future price changes. This suggests that some part of the pollster's surprise is genuine news for the market as well; moreover, it seems to take some time for this new information to be incorporated in the market price.

Regarding the media indices, Table 1 displays evidence of short to medium run causality from price changes to some of these variables. In particular, price changes today are informative about wordcount and printcount in the following day or two; in case of the former possibly even longer. As K increases, these

Table 2: Does $\{x_t\}$ Granger cause $\{p_t\}$? Tests of H_0 : NO vs. H_1 : YES.

	Lags (K)				
Variable	1	2	3	7	10
Election surprise (RCP avg.)	0.000	0.000	0.000	0.000	0.000
Election surprise (latest poll)	0.000	0.000	0.000	0.001	0.006
Wordcount	0.401	0.407	0.064	0.654	0.500
Printcount	0.641	0.626	0.111	0.595	0.731
TV and radio count	0.784	0.177	0.292	0.495	0.317
Webcount	0.724	0.610	0.761	0.676	0.812
Duration	0.952	0.417	0.550	0.434	0.415

Note: The reported figures are p-values.

effects are washed out by the addition of insignificant terms. The rest of the media variables (TV and radio count, webcount and duration) do not seem to respond to price changes conditional on their own history—there is maybe some weak evidence to the contrary for webcount, but none for the other two variables. As shown by Table 2, Granger causality in the other direction (from media indices to price changes) is completely missing. Though the evidence is not overwhelming, these findings suggests that the media coverage of the primary elections consists mostly of noise and stories that are already incorporated in the market price.

To sum up, there seems to be strong two-way Granger causality between prediction market prices and the surprise element in the primaries. There is also some evidence of one-way Granger causality in the short run from price changes towards media news indices. These results suggest that prediction market prices anticipate at least some of the discrepancy between the actual outcome and the latest round of polls before the election. Nevertheless, prices also seem to be *driven* partly by election results, suggesting that there is an element of the pollster's surprise that is genuine news for the market as well.

Table 3: Does $\{p_t\}$ Granger cause $\{x_t\}$? Tests of H_0 : NO vs. H_1 : YES. GARCH effects in price model.

	Lags (K)				
Variable	1	2	3	7	10
Election surprise (RCP avg.)	0.004	0.001	0.000	0.000	0.000
Election surprise (latest poll)	0.323	0.447	0.001	0.000	0.000
Wordcount	0.014	0.031	0.033	0.353	0.002
Printcount	0.000	0.000	0.000	0.000	0.000
TV and radio count	0.008	0.003	0.008	0.001	0.000
Webcount	0.251	0.235	0.265	0.140	0.000
Duration	0.003	0.002	0.001	0.001	0.001

4 Sensitivity analysis and extensions

In order to check the robustness of the baseline results presented in Section 3 we consider some modifications to the model specifications (1) and (2), and sketch possible extensions of the basic linear Granger causality test.

The time series of returns on a financial asset often exhibits conditional heteroskedasticty. Casual visual evidence (the plot of p_t) is consistent with this phenomenon, and more formal tests also show some ARCH effects in the residuals of the models (1) and (2). We therefore repeat the tests described in Section 3 while explicitly modeling conditional heteroskedasticity in the errors as a GARCH(1,1) process. The results are displayed in Tables 3 and 4.

At first glance, explicit modeling of conditional heteroskedasticity changes the numerical results (p-values) quite substantially. Nevertheless, most of the qualitative findings of Section 3 remain true. In particular, there is still fairly strong evidence of mutual Granger causality between price changes and measures of the pollster's surprise, though some of the p-values are now higher (see especially the results in Table 3 concerning surprise relative to the latest polls).

Table 4: Does $\{x_t\}$ Granger cause $\{p_t\}$? Tests of H_0 : NO vs. H_1 : YES. GARCH effects in price model.

	Lags (K)				
Variable	1	2	3	7	10
Election surprise (RCP avg.)	0.007	0.029	0.034	0.002	0.000
Election surprise (latest poll)	0.013	0.054	0.038	0.018	0.076
Wordcount	0.250	0.269	0.304	0.000	0.000
Printcount	0.168	0.093	0.099	0.002	0.000
TV and radio count	0.796	0.396	0.342	0.007	0.035
Webcount	0.870	0.984	1.000	0.014	0.019
Duration	0.691	0.505	0.472	0.000	0.000

Furthermore, as shown by comparing Tables 1 and 3, accounting for conditional heteroskedasticity in price movements seems to amplify the extent to which price changes are informative about future values of the media indices. Printcount and duration, in particular, are now very strongly Granger caused by price movements.

There is however one rather baffling new effect that appears in Table 4. While media variables remain uninformative about future price movements in the short run, it appears that this is not so in the long run (i.e., a week and beyond). This finding is of course very hard to justify theoretically—it not only contradicts the efficient market hypothesis, but also raises the question why there is no short run effect given that there is a long run effect. Further investigation is needed to rule out the possibility that this result is due to a model specification issue.

As discussed in Section 3, the concept of Granger causality used in these tests is rather narrow. More specifically, the results can be interpreted, at best, as testing for causality in the conditional mean. It would be of interest to test separately for causality in mean and causality in variance (or even higher moments) between the price change process and the information variables. One possibility is to use the methodology developed by Cheung and Ng [1996]; this would however necessitate building an explicit univariate model for the information variables. Extending our research in this direction is work in progress.

5 Conclusion

Using data on the race between Barack Obama and Hillary Clinton in the 2008 Democratic primaries for presidential nomination, we investigate whether prediction market prices have information that is not reflected in contemporaneous polls and media stories and conversely, whether prices react to information that appears to be news for pollsters or is prominently featured by the media. We test for Granger causality between day-to-day percent change in prediction market prices and a constructed measure of the surprise element in primary results, i.e information that is not reflected in the polls. We also conduct Granger causality tests between price changes and indices constructed to capture the extent of media coverage received by a candidate.

The main qualitative finding of our exercise, based on the direction of Granger causality found in the data, is that prediction market prices seem to capture some, but not all, of the surprise element in the primary results. Also, there is some evidence that, at least in the short run, prediction market prices capture information that is not reflected in the media. As part of ongoing work, we try to separately test for causality in the conditional variance but this exercise is subject to a number of additional technical difficulties and our preliminary results are sensitive to the exact model specification used.

Appendix

The following is a brief description of the sampling and coding methodology followed by PEJ in constructing their News Coverage Index data library. For further details, see the PEJ website http://www.journalism.org.

The main categories of news sources for the data are as follows: Network TV News, Newspapers, Online News Sites, Cable News and Radio News.

The major broadcast channels ABC, CBS, and NBC make up the broadcast segment. Stories are monitored through different time-slots during the day for 2 out of 3 channels on a rotation basis as follows:

- Commercial Evening News: Entire 30 minutes of 2 out of 3 programs each day (60 minutes)
- Commercial Morning News: 1st 30 minutes of 2 out of 3 programs each day (60 minutes)
- PBS NewsHour: Rotate to code the 1st 30 minutes one day, the 2nd 30 minutes the next day and then skip

This results in either 2 or 2.5 hours of programming each day. Similar method is used on a rotation basis for cable channels CNN, MSNBC and Fox News. During daytime, on a rotation basis, two out of three 30-minute daytime slots each day (60 minutes a day) are coded. During prime time, the following are included

- Two 30-minute segments for Fox News (60 minutes)
- Two 30-minute segments for CNN (60 minutes)
- Two 30-minute segments for MSNBC (60 minutes)

Newspapers are categorized into 3 tiers according to subscription levels, national prominence and regional location. Representative newspapers are chosen from each tier. Here is the list for newspapers in the sample:

- 1st Tier: The New York Times, LA Times, USA Today, Wall Street Journal
- 2nd Tier: Washington Post, Tampa Tribune, Seattle Times, Columbus Dispatch
- 3rd Tier: The Day, Rome News Tribune, Ventura News

For each of the papers selected, only articles that begin on page A1 (including jumps) are picked. This results in a newspaper sample of approximately 20 stories a day.

The websites included in the PEJ sample for our selected date range are as follows: Yahoo News, MSNBC.com, CNN.com, NYTimes.com, Google News, AOL News, Foxnews.com, USAToday.com, Washingtonpost.com, ABCNews.com, HuffingtonPost.com, and Wall Street Journal Online.

For the online news sites, the database captures each site once a day and code the top 5 stories that appear on the site at the time of capture. The time of the day that the sample captures the Web sites is rotated between 9-10 am Eastern time and 4-5 pm Eastern time. The 4-5 PM time-slot was added after April 28, 2008.

The sample of radio stories are collected as follows:

- News: 30 minutes of NPR each day, rotating between Morning Edition and All Things Considered, as broadcast on a selected member station.
- Talk: The first 30 minutes of either one or two talk programs each day. Every weekday, a total of 3 conservatives and 2 liberals were coded during the period of our sample.
- Headlines: Two headline segments each day (one from ABC Radio and

one from CBS Radio), about 10 minutes total.

This results in a sample of roughly 1 or 2 hours of programming a day.

References

K. J. Arrow, R. Forsythe, M. Gorham, R. Hahn, R. Hanson, J. O. Ledyard, S. Levmore, R. Litan, P. Milgrom, and F. D. Nelson. Economics: The promise of prediction markets. *Science*, 320(5878):877, 2008. ISSN 0036-8075.

J. Berg, F. Nelson, and T. Rietz. Accuracy and forecast standard error of prediction markets. *Tippie College of Business Administration, University of Iowa*, 2003.

J. Berg, R. Forsythe, F. Nelson, and T. Rietz. Results from a dozen years of election futures markets research. *Handbook of Experimental Economics Results*, 1:742–751, 2008a. ISSN 1574-0722.

J. E. Berg and T. A. Rietz. Prediction markets as decision support systems. *Information Systems Frontiers*, 5(1):79–93, 2003. ISSN 1387-3326.

J. E. Berg, F. D. Nelson, and T. A. Rietz. Prediction market accuracy in the long run. *International Journal of Forecasting*, 24(2):285–300, 2008b.

K. Y. Chen and C. R. Plott. Information aggregation mechanisms: concept, design and field implementation. *California Institute of Technology Social Science Working Paper*, 1131, 2002.

Y. W. Cheung and L. K. Ng. A causality-in-variance test and its application to financial market prices. *Journal of Econometrics*, 72:33–48, 1996.

R. Engle. New frontiers for arch models. *Journal of Applied Econometrics*, 17(5): 425–446, 2002. ISSN 1099-1255.

R. S. Erikson and C. Wlezien. Are political markets really superior to polls as election predictors? *Public Opinion Quarterly*, 2008. ISSN 0033-362X. doi: 10.1093/poq/nfn010.

J. Geweke, R. Meese, and W. Dent. Comparing alternative tests of causality in temporal systems analytic results and experimental evidence. *Journal of Econometrics*, 21(2):161–194, 1983.

C. W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37:424–438, 1969.

C. A. Sims. Money, income and causality. *American Economic Review*, 62:540–552, 1972.

C. R. Sunstein. Deliberating groups vs. prediction markets (or hayek's challenge to habermas). *Episteme: A Journal of Social Epistemology*, 3(3):192–213, 2007. ISSN 1750-0117.

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

J. Wolfers and E. Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.