International Journal of Forecasting 34 (2018) 696–710

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

Information flow between prediction markets, polls and media: Evidence from the 2008 presidential primaries

Urmee Khan [a,*], Robert P. Lieli [b]

[a] Department of Economics, University of California, Riverside, United States
[b] Department of Economics, Central European University, Budapest, Hungary

ARTICLE INFO

Keywords:
Prediction markets
Opinion polls
Information aggregation
Media coverage
Granger causality

ABSTRACT

Are the forecast errors of election-eve polls themselves forecastable? We present evidence from the 2008 Democratic Party nomination race between Barack Obama and Hillary Clinton showing that the answer is yes. Both cross-sectional and time series evidence suggests that market prices contain information about election outcomes that polls taken shortly before the contests do not. Conversely, election surprises relative to polls too Granger cause subsequent price movements. We then investigate whether the additional information in prices could come from the media coverage of these campaigns, and uncover a set of complex relationships between pollster's surprise, price movements, and various aspects of media coverage. Prices anticipate the *balance* and *content* of media coverage, but not the *volume*. On the other hand, it is the volume of media coverage, not the balance or content, that anticipates the surprise element in election outcomes. Moreover, Granger causality between prices and election surprises barely changes after controlling for media coverage, and causality from media volume to surprises persists too after controlling for price movements. Taken together, the results suggest that both prices and the volume of media coverage contain independent election-relevant information that is not captured in polls.

1. Introduction

The efficacy of political forecasting has been front and center in public discourse following recent U.S. presidential election cycles. The widespread failure of the forecasting industry to identify Donald Trump as the likely winner in 2016 has set off an intense scrutiny of popular forecasting methodologies, much of which has focused on public opinion polls and the possible causes of mismeasurement (Bialik & Enten, 2016; Lohr & Singer, 2016; Mercer, Deane, & McGeeney, 2016). Similar questions were raised after the 2008 primary cycle, in light of the fact that the eventual winner of the Democratic Party primaries in many competitive states was consistently underpolled (Traugott et al., 2008; Traugott & Wlezien, 2009). While the failure of prediction markets has received less public attention, the fact remains that they, too, fared very poorly in forecasting the outcome of the last presidential election. A rather stunning illustration is provided by the daily evolution of vote share and winner-take-all security prices in the Iowa Electronic Markets (IEM). Even a casual glance at these time series shows that markets considered Hillary Clinton as the overwhelming favorite up until the very last moment.[1]

The use of pre-election *voter intention* polling as an election forecasting tool is one of the most widely studied topics in public opinion research, and the literature is vast (see for example Arzheimer, Evans, & Lewis-Beck, 2017, Ch. 36, for an overview). From a statistical point of view,

* Corresponding author.
E-mail addresses: urmeek@ucr.edu (U. Khan), lielir@ceu.edu (R.P. Lieli).

[1] See http://tippie.biz.uiowa.edu/iem/. When accessed on Sept. 18, 2017, these figures were still displayed prominently on the front page.

https://doi.org/10.1016/j.ijforecast.2018.04.002
0169-2070/

one can expect intention polls to be unbiased and accurate forecasts of election outcomes only if several conditions are met: (i) the respondents are a representative sample of the subset of the population that turns out to vote; (ii) the respondents reveal their intentions truthfully, and (iii) the respondents do not change their intentions systematically before the election date. In practice all of these conditions are likely to be violated to some degree (see Gass & Fu, 2013, pp. 539–604, for a discussion of these issues), and even if they are met, random sampling variation still endows polls with a margin of error.

As an alternative to voter intentions, the election forecasting literature has also proposed polling *voter expectations* of the election outcome (known as 'citizen forecasts'; see for example Lewis-Beck & Skalaban, 1989; Lewis-Beck & Stegmaier, 2011; Lewis-Beck & Tien, 1999; and Murr, 2011, 2015, 2016). As Graefe (2014) and Rothschild and Wolfers (2013) show, expectation polls may indeed outperform intention polls for forecasting election outcomes and vote shares. However, such polls are still rare in practice, despite the theoretical appeal. Prediction markets, on the other hand, have become ubiquitous, and the theoretical and empirical literature on the role of these markets in political forecasting has grown rapidly over the last two decades (Arrow et al., 2008, and Chen & Plott, 2002, were early endorsements; see e.g. Horn, Ivens, Ohneberg, & Brem, 2014, Tziralis & Tatsiopoulos, 2007, and the references therein for a broad overview).

Prediction markets are similar to polling-based citizen forecasting approaches, in that they elicit information broader in scope than just one's own intent. However, they do so in an incentivized way, and from a group of self-selected, rather than randomly selected, participants. On the other hand, intention polls elicit only one's own intent, are not incentivized, and attempt to capture the opinions of the voting population 'at large'. Based on these differences, comparative assessments of polls and prediction markets as information gathering tools are important, and such an exercise is the main objective of the paper. The exercise is motivated by, and related to, multiple strands of literature that study polls and prediction markets either individually or in a comparative setting.

In comparing voter intention polls to prediction markets, the empirical literature has focused typically on forecast accuracy metrics (primarily the mean squared error). In a series of papers, Berg, Nelson, and Rietz (2003, 2008) use IEM data to argue that prediction market prices have significantly lower forecast errors than contemporaneous polls in both the short and long run. Results in a similar vein are obtained by Leigh and Wolfers (2006). However, other concurrent research paints a more complicated picture. Jacobsen, Potters, Schram, van Winden, and Wit (2000) present evidence of market-based forecasts for European elections not being particularly accurate. Using data from the same time period as Berg et al. (2008), Erikson and Wlezien (2008) construct projections of vote shares and win probabilities based on daily polls and show that these projections can outperform prediction markets. Using a long series of historical data on election betting markets dating back to the early 20th century, Erikson and Wlezien (2012a) present a comparative analysis of markets before and after the introduction of scientific polling, and provide intriguing evidence that market prices were far better predictors in the period without polls than in the period once polls became available. Finally, as was noted above, the 2016 U.S. election cycle provided a cautionary tale that applies to markets as much as to polls.

Given that the empirical evidence regarding the relative superiority of either intention polls or markets as standalone forecasts is hardly decisive, it is natural to use either or both as inputs to more sophisticated, combined forecasting methods. Prominent examples range from the classic political economy models pioneered by Lewis-Beck and Rice (1982) (see also Arzheimer et al., 2017, Ch. 26, and Lewis-Beck & Rice, 1984; Lewis-Beck & Tien, 1996, gives an overview of recent research) to the likes of the 'polls-plus' forecast from FiveThirtyEight.com and other multi-input political forecasts from websites such as PollyVote.com and PredictWise.com.

The mixed evidence on the comparative forecasting performances of intention polls and markets in various settings, coupled with the improvements that are seemingly afforded by combined forecasts, points to possible differences in their information content. The goal of this paper is to investigate the 'information gap' between intention polls and other information aggregation mechanisms *explicitly*, rather than indirectly through forecast accuracy metrics. The approach that we take is based on a simple idea. A well-known property of optimal forecasts under square loss is that the associated forecast errors must be uncorrelated with any variable in the forecaster's information set; more precisely, the forecast errors themselves must not be forecastable using any function of these variables, though they may still be forecastable using variables from a larger information set.

Motivated by these observations, the first question that we ask is whether the discrepancy between the actual election outcomes and the last round of polls, treated as a short-run forecast error, is anticipated by other information aggregation mechanisms, and in particular by preceding market prices. In the reverse direction, the forecast errors made by polls may also predict post-election market price movements if they are correlated with the forecast errors made by the market (i.e., if the polling error arose for reasons that the market also failed to account for), and there is also some degree of persistence in the adjustment of prices.

While this exercise by itself adds to the previously cited literature on poll-market comparisons, it does not reveal what the market may potentially know in addition to poll numbers, or whether there might be other types of public information that are predictive of polling errors but are not incorporated into market prices.[2] The primary source of public information about political campaigns is media coverage. This includes coverage of poll numbers, but potentially also other information that (a) might help to adjust for systematic biases in polls, or (b) is slow to

[2] This last scenario implies an apparent failure of the efficient market hypothesis when applied to prediction markets. This is often assumed directly in related theoretical work (e.g., Kou & Sobel, 2004), but we think that it is best regarded as an empirical question.