THE 2016 POLLYVOTE POPULAR VOTE FORECAST:

A PRELIMINARY ANALYSIS

Andreas Graefe, J. Scott Armstrong, Randall J. Jones, Jr., and Alfred G. Cuzán

Since its launch last January, the 2016 PollyVote consistently predicted that Hillary Clinton would win the popular vote, which she did. In this preliminary analysis we assess how the PollyVote and its components performed in this election compared to the previous six (1992 to 2012).

Ranked according to their historical accuracy from best to worst, right to left in Figure 1, the six compoments of the PollyVote place as follows: citizen forecasts, prediction markets, index models, expert judgment, econometric models, and polls. In 2016 citizen forecasts[1] defended their top position, but other methods were outliers relative to their historical record. Both polls and econometric models performed considerably better, while prediction markets did worse. This fluctuation in accuracy among methods makes it difficult to predict *ex ante* which method will do best.

Combining forecasts within and across methods, as the PollyVote does, guarantees that its forecast will never come out last and will always do at least as well as the typical forecast.[2] This was demonstrated this year. The PollyVote incurred an error of 1.9 percentage points, almost twice the historical average, but lower than that of three other methods.

. . . .

Figure 1 about here

. . . .

Combining works best when the various component forecasts bracket the true value (Graefe et al., 2014a). Compared to the previous six elections, there wasn't

much bracketing in 2016. Five of the combined forecasts overshot Clinton's share of the vote while only one component, the econometric models, fell short. Thus the PollyVote did not do as well as in previous elections.

The method of combining forecasts makes no claim that the PollyVote will always outperform its most accurate component, although that can happen, as was the case in 2004 (Cuzán et al., 2005) and 2012 (Graefe et al., 2014b). What *is* claimed is that over time, as the constituent methods' relative accuracy varies, the PollyVote will surpass them. This is demonstrated in Figure 2, which displays the mean absolute error of all methods across all seven elections from 1992 to 2016. On average, then, the PollyVote continues to minimize error while avoiding making large errors.

. . . .

Figure 2 about here

. . . .

REFERENCES

Cuzán, A. G., Armstrong, J. S. & Jones, R. J. Jr. (2005). How we computed the PollyVote. *Foresight: The International Journal of Applied Forecasting*, 1(1), 51-52.

Graefe, A. (2014). Accuracy of Vote Expectation Surveys in Forecasting Elections, *Public Opinion Quarterly*, 78, S1, 204-232.

Graefe, A., Armstrong, J. S., Jones, R. J. Jr. & Cuzán, A. G. (2014a). Combining forecasts: An application to elections, *International Journal of Forecasting,* 30(1), 43-54.

Graefe, A., Armstrong, J. S., Cuzán, A. G. & Jones, R. J. Jr. (2014b). Accuracy of combined forecasts for the 2012 presidential elections: The PollyVote, *PS: Political Science & Politics,* 47(2), 427-431.

Figure 1. Forecast error by method
(Mean absolute error, historical vs. 2016, across last 100 days before the election)

Historical   2016
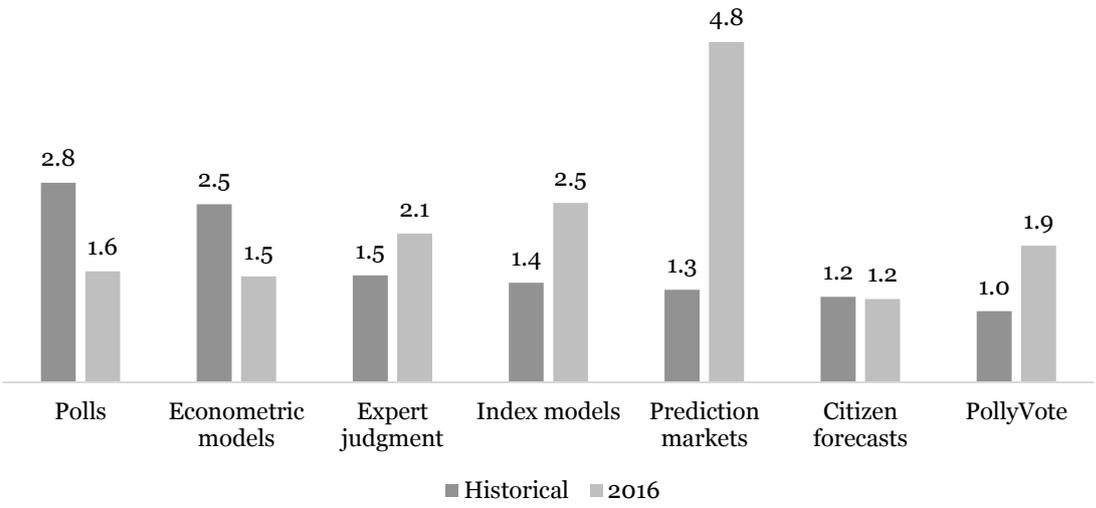
Figure 2. Forecast error by method

(Mean absolute error, 1992-2016, across last 100 days before the election)
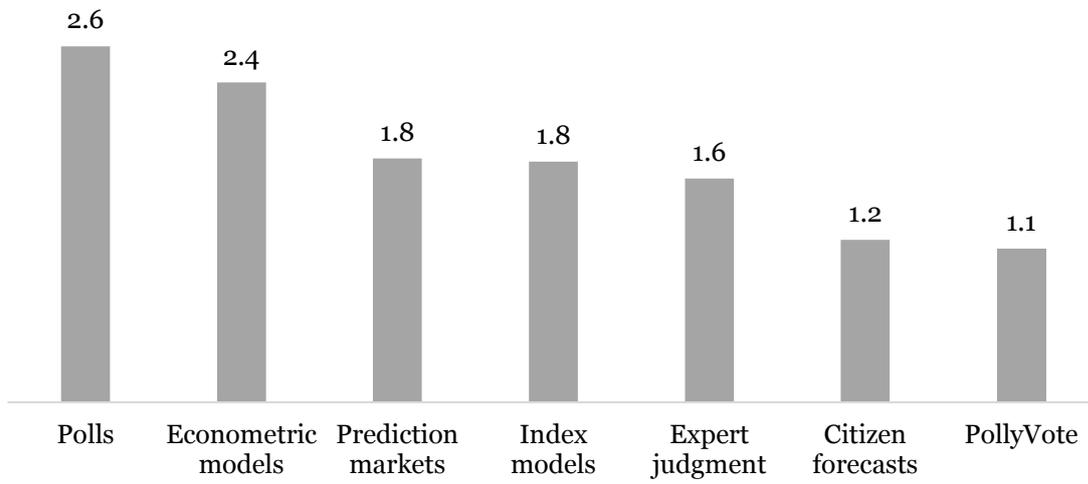
¹ On the remarkable performance of citizen forecasts, see Graefe (2014).

² The error of the typical forecast is *the average of the errors* incurred by individual forecasts of a given event, such as an election. By contrast, the error of the combined forecast is the difference between *the average of the forecasts* and the true value of the event being predicted.