

**Predicting elections from the most important issue:
A test of the take-the-best heuristic**

Forthcoming in the *Journal of Behavioral Decision Making*

Andreas Graefe
Institute for Technology Assessment and Systems Analysis
Karlsruhe Institute of Technology, Germany
graefe@kit.edu

J. Scott Armstrong
The Wharton School
University of Pennsylvania, Philadelphia, PA
armstrong@wharton.upenn.edu

Abstract. We used the take-the-best heuristic to develop a model to forecast the popular two-party vote shares in U.S. presidential elections. The model draws upon information about how voters expect the candidates to deal with the most important issue facing the country. We used cross-validation to calculate a total of 1,000 out-of-sample forecasts, one for each of the last 100 days of the ten U.S. presidential elections from 1972 to 2008. Ninety-seven percent of forecasts correctly predicted the winner of the popular vote. The model forecasts were competitive compared to forecasts from methods that incorporate substantially more information (e.g., econometric models and the Iowa Electronic Markets). The purpose of the model is to provide fast advice on which issues candidates should stress in their campaign.

The common citizen generally has little interest in political matters. Many people do not know how government works or have opinions on the major issues facing the county. Nonetheless, Lau and Redlawsk (1997) found that about 75% of voters in the five U.S. Presidential Elections from 1972 to 1988 voted “correctly”. That is, people voted for the candidate that they would have also chosen if they had possessed full information. The new conventional wisdom in political science is that voters overcome their cognitive limits by using simple decision rules, commonly referred to as heuristics. Heuristics enable voters to make reasonable decisions without possessing complete information about candidates and issues. For example, instead of evaluating candidates on multiple attributes, voters might compare candidates only on a limited set of attributes. A common heuristic is *single-issue voting* whereby voters compare candidates only on the issue that is most important to them.

Several conditions are expected to impact the use of heuristics. For example, people can be expected to rely on heuristics more often if the decision is not very important. In the case of U.S. presidential elections, the impact of a single vote on the election outcome is marginal. The decision to vote might be motivated rather by fulfilling a civic duty than to find the best option. In addition, heuristic use can be expected to be high in situations where people have to process a lot of information and comparison of alternatives is difficult. Redlawsk (2004) reported experimental results showing that heuristic use increases in more complex environments. In the information rich environment of U.S. presidential elections, it seems impractical for voters to evaluate the candidates on their positions on the complete set of issues. And how can voters evaluate candidates on incomparable attributes? For example, voters might have difficulties in weighing the candidates’ stands on the issues of abortion and the war in Iraq.

If a single issue dominates a campaign (such as the state of the economy in the 2008 U.S. Presidential Election), many voters may make their voting decision solely based on the candidates’ ability to handle this issue. We used the *take-the-best heuristic* to develop a model for predicting the popular two-party vote in U.S. Presidential Elections that uses information about which candidate is favored by voters for dealing with the most important issue. The goal was to develop a model that can provide fast advice on which issues candidates should stress in their campaign. Our approach uses publicly available polling data about perceptions of issue importance and the candidates’ ability to handle the one issue seen as most important by voters.

Take-the-best heuristic

Take-the-best (TTB), developed by Gigerenzer and Goldstein (1996), is a heuristic for choosing between alternatives based on a single piece of information. That is, starting with the most important attribute, a decision maker looks whether this attribute discriminates between the alternatives. If yes, he makes a decision for the alternative that is favored by the attribute. If not, he moves on to the next most important attribute.

Czerlinksi et al. (1999) found TTB to be remarkably predictive. The authors compared the heuristic to multiple regression and unit-weighting for 20 prediction problems (e.g. forecasting high school drop out rates, attractiveness of men and women, homelessness and mortality in U.S. cities, salaries of college professors, obesity among children, and fish fertility) for which the number of variables varied between 3 and 19. Most of these examples were taken from statistical textbooks where they had the purpose of demonstrating the application of multiple regression analysis. Not surprisingly, multiple regression performed best when calculating in-sample forecasts. However, when using cross-validation to predict data that had not been used to build the model, TTB was most accurate, followed by unit-weighting. The advantage of TTB was higher when there were fewer observations per predictor variable. Even for more than ten observations per variable, multiple regression seldom outperformed TTB.

Martignon and Hoffrage (1999) discovered further conditions for the applicability of TTB. They showed mathematically that a linear model cannot outperform TTB if the implicit

importance weight of a variable is greater than the sum of the weights of all less important variables.

TTB seems to provide a promising way to predict U.S. presidential election outcomes. First, the number of observations is small. For forecasting U.S. presidential elections, data for the majority of regression models is limited to about 25 elections. Information about how voters perceive candidates to handle the issues was available only for the last ten elections. Second, a single most important issue often dominates presidential campaigns. Third, it is now easy to obtain information about the importance of issues. Polling institutions frequently ask voters to name the single most important issue facing the country as well as to state which candidate can do a better job in handling a particular issue. In recent years, the Internet has made such information more readily available

Big-issue voting model

We developed the *big-issue* (BI) *voting* model to predict the outcome of U.S. presidential elections. The model is based on information about how voters expect the candidates to deal with the issue seen as most important. The model relies solely on information from polls and uses a heuristic similar to TTB to determine the winner of the popular vote.

Data

We collected polling data on which issue voters regard as being the single most important issue facing the country for the period from June to October in each election year (e.g., “*What do you think is the most important problem facing this country today?*”, Gallup Poll, October 3-October 5, 2008). To obtain this, we searched the *iPOLL databank* by the *Roper Center for Public Opinion Research*, using the search string “most important issue OR most important problem”. (Note that *issue* and *problem* are used interchangeably. See Wlezien (2005) for a different view.)

We obtained polling data for the last ten U.S. Presidential Elections from 1972 to 2008. For example, during the 2008 election, voters consistently saw the state of the economy as the most important issue. By comparison, in 2004 the economy and the war in Iraq alternated as the most important issue.

We also obtained polling data on which candidate voters believed would be more successful in solving the issue seen as most important. Again, we searched the *iPOLL databank* by using the following search string: “[Republican candidate] AND [Democratic candidate] AND [most important issue]”. For example, for the 2008 election, we searched for “McCain AND Obama AND economy” and analyzed poll questions such as: “*Regardless of which (2008) presidential (election) candidate you support, please tell me if you think Barack Obama or John McCain would better handle each of the following issues. How about...the economy?*” (cf. Gallup/USA Today Poll, October 10-October 12, 2008).

BI heuristic to determine the election winner

We used a variant of TTB to predict the winner of the popular vote in an election campaign. We refer to this approach as the *BI heuristic* (BI-H). The underlying decision rule performs three basic steps: (1) Identify the issue seen as most important by voters, (2) calculate the two-party shares of voter support for the candidates on this issue and average them for the last three days, and (3) predict the candidate with the higher voter support to win the popular vote. For a more detailed description of our decision rule, see Appendix 1.

Note that the heuristic does not require information about historical elections to determine the election winner. All necessary information can be derived from the polls in the respective election year.

BI model to predict the two-party vote shares

We used a simple linear regression to build the *BI model* (BI-M). We related the incumbent party candidate's 3-day rolling average of voter support on the most important issue (referred to as S below) to the dependent variable, which was the actual two-party popular vote share received by the candidate of the incumbent party (V). Based on data from the last forecasts prior to Election Day in the ten elections from 1972 to 2008, we derived the following vote equation:

$$V = 27.0 + 0.50 * S$$

Thus, BI-M predicts that an incumbent would start with 27% of the vote, plus a share depending on S. If the incumbent's voter support on the most important issue went up by 10 percentage points, the incumbent's vote share would go up by 5 percentage points. Consistent with traditional econometric models, BI-M reveals an advantage for the incumbent. Assuming the candidates achieve equal voter support (i.e., S=50.0), the candidate of the incumbent party would be predicted as the winner, receiving a vote share of about 52%.

Forecasting accuracy of BI

For each election year, the forecast origin was 100 days prior to Election Day, which was moved forward by one day at a time until Election Eve. Thus, over the ten U.S. Presidential Elections from 1972 to 2008, we obtained a sample of 1,000 forecasts.

Predicting the election winner

The performance of BI for predicting the winner (i.e., the hit rate) varied during the election campaign as new polls became available. The hit rate is the proportion of forecasts that correctly determined the winner of the popular vote.

Over all 1,000 forecasts, BI-H correctly predicted the winner 88% of the times. This performance was achieved without incorporating information from previous election years.

BI-M was more accurate. The model's forecasts were calculated through N-1 cross-validation. This means that we used the observations from 9 elections to calibrate the model and then made a forecast for the one remaining election. Over all 1,000 forecasts, BI-M correctly predicted the winner 97% of the times. The model's hit rate was especially accurate close to Election Day: on each of the last five days prior to Election Day, the BI-M forecast correctly predicted the winner in each of the ten elections.

Combining forecasts of BI-H and BI-M was expected to further increase accuracy. BI-H predicts the incumbent to win only if the incumbent achieved a higher voter support than the opponent. In incorporating an advantage for the incumbent, BI-M can still predict the incumbent to win, even if the opponent has a slight advantage in voter support for the most important issue.

Table 1 shows combined forecasts of BI-H and BI-M. The combined forecasts reveal how often both approaches agreed in their forecasts of the winner. In 87% of all forecasts, both approaches correctly predicted the winner, whereas they simultaneously missed the winner only 2% of the times. The cases in which the forecasts from both approaches were wrong were limited to the elections in 2000 and 2004, both of which were very close. In 11% of the cases, BI-H and BI-M disagreed, which can be useful as an indicator of uncertainty.

——— Table 1 about here ———

Predicting the incumbent's two-party share of the popular vote

Although it is not the primary goal of this study to improve the accuracy of election forecasts, we used well-established election forecasting methods as a benchmark to assess accuracy. Most models issue their forecast around Labor Day in the respective election year, which is usually

about eight to nine weeks prior to Election Day. Table 2 reports the forecasts of BI-M calculated 9 weeks, or 63 days, before Election Day. Again, the forecasts were calculated through N-1 cross validation. Over all 10 elections, the mean absolute error (MAE) was 3.0 percentage points.

——— Table 2 about here ———

The critical test is how well the models forecast prospectively (that is, for years not included in the estimation sample). We generated *ex ante* forecasts for the last three elections from 2000 to 2008 by successive updating. That is, only data from elections prior to the respective election year were used for building the model. To predict the 2008 election, data on the 9 elections from 1972 to 2004 were used; for the 2004 election, data on the 8 elections from 1972 to 2000 were used, and, for the 2000 election, data on the 7 elections from 1972 to 1996 were used.

Table 3 shows *ex ante* forecasts from BI-M and eight econometric models. Most of the forecasts from the econometric models were published in *PS: Political Science and Politics*, 34(1), 37(4), and 41(4). The forecasts for Fair's model were obtained from his website (<http://fairmodel.econ.yale.edu>). For an overview of the predictor variables used in most of the models, see Jones and Cuzán (2008).

——— Table 3 about here ———

Although drawing on less information, the BI-M provided competitive forecasts. It yielded a lower MAE than four of the eight econometric models and than the typical model. Furthermore, one model that yielded a lower MAE than BI-M predicted the wrong winner in the 2004 election. BI-M correctly predicted the winner in all three elections.

Relative accuracy of BI and the Iowa Electronic Markets

In recent elections, the number of polls has grown rapidly and new polls are published almost on a daily basis. As a result, the BI forecasts may change frequently. Thus, we compared the daily forecasts of BI-H and BI-M to forecasts from prediction markets, which also provide daily updated forecasts.

Betting markets to predict election outcomes have an interesting history. Rhode and Strumpf (2004) studied historical betting markets that existed for the 15 presidential elections from 1884 through 1940 and concluded that these markets “did a remarkable job forecasting elections in an era before scientific polling” (2004:127). In 1988, the *Iowa Electronic Market* (IEM) was launched as a futures market in which contracts were traded on the outcome of the presidential election that year. Initially, the IEM, commonly viewed as a *prediction market*, provided more accurate election forecasts than trial-heat polls. In comparing IEM vote-share prices with 964 trial-heat polls for the five presidential elections from 1988 to 2004, Berg et al. (2008) found that IEM market forecasts were closer to the actual election results 74% of the time.

We compared the relative performance of BI and the IEM's *vote-share* and *winner-take-all* markets. The vote-share markets provide a quantitative forecast of the two-party popular vote-shares achieved by the candidates. Winner-take-all markets provide a forecast of which candidate will win the popular vote. For the six elections from 1988 to 2008, Table 4 shows the hit rates of BI-H, BI-M (forecasts calculated by N-1 cross validation), and the two IEM markets for the last 100 days prior the Election Day. For each day, we used the last traded price of the two IEM markets.

——— Table 4 about here ———

BI-H and BI-M performed well and yielded higher hit rates than the vote-share markets. The winner-take-all markets, available from 1992 to 2008, performed as well as BI-H, but were inferior to BI-M, which achieved 94% correct predictions.

By comparison, the IEM vote-share markets were more accurate than BI-M in predicting the actual vote-shares. Across the 100 days in the forecasting horizon, the MAE over the six elections from 1988 to 2008 was 1.7 for the vote-share markets and 2.4 for BI-M. The superiority of the IEM for predicting vote shares is not surprising as participants in prediction markets can draw on information from many sources, most notably trial-heat polls. However, the results suggest that BI-M may improve on the accuracy of prediction markets by providing an accurate prediction of the election winner.

Discussion

The state of election forecasting has progressed to the point where it is possible to get highly accurate forecasts for major elections. However, one area that has received little attention is how to use forecasting as an aid to those involved with political campaigns. Although traditional econometric models reveal that the incumbent's chances to win decrease if he is unpopular, the economy is doing poorly, or the federal budget deficit increased during his administration, it is difficult for political parties and candidates to take action based on such forecasts. Such models provide limited or no advice with respect to questions such as what type of candidate a party should nominate or what issues should be stressed in the campaign.

Implications of the BI model for campaign strategies

In using information about how voters perceive the candidates to handle the most important issue, the BI model can be used for advising those involved in political campaigns. To the best of our knowledge, the BI model is the first model that incorporates such information.

The BI model uses TTB to illustrate how information on issues can be effective for forecasting the election outcome. The model does not aim at predicting individual voting behavior. While individual voters may use many different voting strategies, the model suggests that the electorate as a whole behaves as if its voting decision is based on the candidates' ability to handle the single most important issue.

The BI model is simple to use and easy to understand, and it has decision-making implications. Political candidates can take action upon the forecast and develop campaign strategies. Two different strategies seem conceivable for allocating campaign efforts: (1) to be seen as the best candidate for dealing with the most important issue or (2) to change voters' perceptions about which issue is most important.

The first strategy seems difficult to pursue as voter perception of which candidate will do better in handling a particular issue seldom changes. The reason is party identification. Each party has a distinct issue handling reputation, depending on their party constituency, which changes very slowly if at all. In reporting a small sample of polling data for the period from 1988 to 1991, Petrocik (1996) found that Democrats were favored for welfare issues (e.g. health care and social security), whereas Republicans had advantages for social issues (e.g. crime and upholding moral values) and issues related to foreign policy and defense. Perceptions of the parties' performance on handling economic issues were mixed.

The second strategy seems more promising; that is, trying to change voters' perceptions about the importance of issues. An example is the 2004 election: although only two issues were seen as most important (the war in Iraq and the economy), voters' perceptions about the importance of these issues changed seven times, more than in all other elections. While Bush was favored on solving the issue with the war in Iraq, Kerry was seen as more competent in dealing with the economy. During the 100 days prior to Election Day, the economy was seen to be the most important issue on 86 days, despite the fact that foreign policy (i.e. the situation in Iraq and the war against terrorism) was the dominant theme throughout the campaign. However, in the six polls during the last five days prior to Election Day, the war in Iraq was on average seen as most important. Our model suggests that this change in issue importance helped Bush win the election. Similarly, in the 2000 election, our model incorrectly predicted Bush to win the popular vote

early in the campaign, when moral values were seen as the most important issue. Bush's early lead might have been a response to his attacks of the Clinton scandals: restoring "honor and dignity" in the White House was a major theme in Bush's campaign. However, later in the campaign, most likely due to effective campaigning by Gore, domestic issues like health care, social security, and education caught the attention of the public.

These observations suggest that candidates should aim at promoting issues for which they are favored in the eyes of voters. In fact, candidates have often been found to behave this way. In analyzing *New York Times* articles, published in October in the election years from 1952 to 1988, Petrocik (1996) found that candidates focused on issues for which voters traditionally favored their party and rarely attempted to change voters' opinion on issues that favored the opposing party.

Spotting emerging issues

Since issue-handling reputations are difficult to change, candidates are advised to gain competence on issues that are likely to emerge. Thus, spotting emerging issues becomes an important task for political candidates. A new approach that can be useful for the early identification of emerging issues relies on information about Internet searches. Search engines provide information about how many searches have been done over time for certain keywords in a certain geographical area (relative to the total number of searches or to the number of searches within a predefined category). Thus, search query data should be useful to measure the importance of certain issues: the more users search for a certain issues, the greater is the importance that is likely to be attached to it.

We obtained search query data from *Google Insights for Search* to analyze how the importance of the two most important issues in the 2004 election changed in the run-up to the 2008 election. Figure 1 shows the growth in searches (with respect to the first date on the graph) for the keywords "Iraq" and "economy" from January 2004 to November 2008. (The data is limited to searches that were conducted within the U.S. and within the category "Society".) As can be seen, searches for "Iraq" were most popular during most of 2004. However, since September 2007 searches on "economy" took over. Even though the war in Iraq was still a key issue in the 2008 campaign before the economic crisis, keyword searches already indicated the increasing importance of economic issues.

– Figure 1: Keyword popularity on Google Insights for Search for "Economy" and "Iraq" (2004 - 2008) –

The BI model provides an easy and intuitive way to track campaigns. This is a major advantage of the model compared to traditional econometric models. Our results suggest that candidates should focus on raising and promoting issues for which they are favored among voters, especially new issues that have not yet received much public attention.

Conditions for the BI model

The BI model is expected to work well if the most important issue is of widespread interest and intense public importance. However, it may lead to poor forecasts if there is no single issue that is clearly more important than others. In such situations, issues that were ignored may yield a different preference order among candidates.

The data appeared to conform to these expectations. Table 5 lists the performance of the BI combined forecasts in predicting the election winner, along with the issues seen as most important by voters in each election year. The BI combined forecasts were particularly accurate in elections, where (i) the most important issue never changed (as in 1972, 1976, and 2008) or (ii) the issues seen as most important belonged to the same category (as for economic issues in 1980 and 1992). If a candidate is favored for handling the issue of inflation, it is likely that he is also favored for handling the economy in general.

——— Table 5 about here ———

By comparison, the BI combined was least accurate for the three elections in 1988, 2000, and 2004. For these elections, voter perception of which issue is most important changed most frequently and/or the most important issue belonged to different categories. In such situations, focusing on the single most important issue might not be enough. It might be worthwhile to consider information on a larger number of issues.

Issue-indexes: An alternative to the single most important issue

Given a situation under which (i) a large number of causal variables are important and (ii) there is good prior knowledge on the importance of variables (and the directions of effects), one might want to consider the *index method*. In using this old forecasting method, analysts prepare a list of key variables and specify from prior evidence whether they are favorable (1) or unfavorable (0) in their influence on a certain outcome. Then, the analysts simply tally the scores and use the totals to calculate the forecast. The index method also draws on the idea of *unit-weighting*, also known as *Dawes' rule* or *tallying*.

Lichtman (2008) was the first to use the index method for election forecasting. His “Keys” model for forecasting U.S. presidential election winners assigns values of zero or one to an index of thirteen predictor variables. The model predicts the incumbent party to lose the popular vote, if it loses six or more of the thirteen keys. Although the keys include two measures of economic conditions, they do not include information on specific issues that concern voters in a particular election. Instead, the keys include questions such as whether the incumbent president was involved in a major scandal or whether the current administration was affected by foreign or military success (or failure). The “Keys” model provided correct forecasts retrospectively for all of 31 elections and prospectively for all of the last 7 elections. No model has matched this level of accuracy in picking the winner of the popular vote.

Armstrong and Graefe (in press) used an index of 59 biographical variables to predict the popular vote winner in the 29 U.S. presidential election winners from 1896 to 2008. The variables measured, for example, whether a candidate was married, went to a prestigious college, or was taller than the opponent. This biographical index model correctly predicted the winner in 27 of the 29 elections and yielded out-of-sample forecasts that were as accurate as the best of seven econometric models.

In addition to providing accurate forecasts, both models have decision-making implications for political campaigns as they advise political parties in who they should nominate. For example, the “Keys” model includes variables that suggest that parties should choose candidates who are considered national heroes or highly charismatic. The biographical index model provides an extensive checklist that helps parties to evaluate the chances of different candidates to win an election.

The index method is an effective approach for including much information in a model. Thus, it should be useful for capturing information about how voters perceive the candidates' ability to handle a set of issues. Such information could be a valuable addition to the BI model as it could increase certainty in the forecast; *in particular in election campaigns in which many issues are seen as important*. We are currently investigating whether an issue-index model might offer additional useful information for election campaigns (Graefe & Armstrong 2010).

Predicting multi-party elections

Although the BI model was developed and tested for the two-party system of U.S. presidential elections, it may also be valuable for forecasting election outcomes in multi-party systems. However, generalizing TTB to multiple-alternative outcome tasks is not without difficulties.

For example, the decision-maker has to decide how to proceed if there is no one best alternative for the most important attribute. Rieskamp and Hoffrage (1999) suggested that, when moving on to the next most important attribute, one could either compare against all alternatives or compare only the alternatives that tied on the previous attribute.

Similar problems occur if there is no single most important attribute. Imagine a situation with three parties (A, B, C) and two issues (x, y) of similar importance. Voter support for party A is 55% for issue x and 10% for issue y; or A(55,10). The respective numbers for parties B and C are B(10,55) and C(35,35). If a decision-maker makes a decision solely based on issue x (or y), he would choose party A (or B). If he would average voter support across issues, he would choose party C. Note that our decision rule, described in Appendix 1, avoids this problem. If necessary, the decision rule averages results from previous polls and thus ensures that it is always possible to identify the most important issue. Future research should evaluate the usefulness of the BI model for predicting multi-party elections.

Conclusions

The take-the-best heuristic generated accurate forecasts based on voters' perceptions on how the candidates will handle the single most important issue facing the country. In a cross-validation of 1,000 out-of-sample forecasts for the ten U.S. presidential elections from 1972 to 2008 (one forecast on each of the last 100 days per election year), the model correctly predicted the winner of the popular vote in 97% of all forecasts. This quick and simple model can help candidates to develop campaign strategies.

For the six elections from 1988 to 2008, the model yielded a higher number of correct predictions of the popular vote winner than the Iowa Electronic Markets. For the three elections from 2000 to 2008, the BI model forecasts of the incumbent's two-party share of the popular vote yielded accuracy similar to the typical econometric model. In using a different method and different information than traditional approaches, the BI model is expected to contribute to forecasting accuracy and should be combined with other forecasting models.

Acknowledgments

Andrew Gelman, Dan Goldstein, and three anonymous reviewers provided helpful comments. Kesten C. Green provided peer review.

References

- Armstrong, J. S. & Graefe, A. (in press). Predicting elections from biographical information about candidates: A test of the index method. *Journal of Business Research* (forthcoming).
- Berg, J. E., Nelson, F. D. & Rietz, T. A. (2008). Prediction market accuracy in the long run. *International Journal of Forecasting*, 24, 285-300. DOI:10.1016/j.ijforecast.2008.03.007
- Czerlinski, J., Gigerenzer, G. & Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer, P. M. Todd & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 97-118). Oxford University Press.
- Gigerenzer, G. & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103, 650-669. DOI: 10.1037/0033-295X.103.4.650
- Graefe, A. & Armstrong, J. S. (2010). *Forecasting elections from voters' perceptions of candidates' ability to handle issues*. Available at SSRN: ssrn.com/abstract=1206683.
- Jones, R. J. & Cuzán, A. G. (2008). Forecasting U.S. presidential elections: A brief review. *Foresight – The International Journal of Applied Forecasting*, Issue 10, 29-34.
- Lau, R. R. & Redlawsk, D. P. (1997). Voting correctly. *American Political Science Review*, 91, 585-598.
- Lichtman, A. J. (2008). The keys to the white house: An index forecast for 2008. *International Journal of Forecasting*, 24, 299-307. DOI: 10.1016/j.ijforecast.2008.02.004

- Martignon, L. & Hoffrage, U. (1999). Why does one-reason decision making work? A case study in ecological rationality. In G. Gigerenzer, P. M. Todd & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 119-140). Oxford University Press.
- Petrocik, J. R. (1996). Issue ownership in presidential elections, with a 1980 case study. *American Journal of Political Science*, 40, 825-850.
- Redlawsk, D. P. (2004). What voters do: Information search during election campaigns. *Political Psychology*, 25, 595-610. DOI: 10.1111/j.1467-9221.2004.00389.x
- Rieskamp, J. & Hoffrage, U. (1999). When do people use simple heuristics, and how can we tell? In G. Gigerenzer, P. M. Todd & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 141-167). Oxford University Press.
- Rhode, P. W. & Strumpf, K. S. (2004). Historical presidential betting markets. *Journal of Economic Perspectives*, 18, 127-141. DOI:10.1257/0895330041371277
- Wlezien, C. (2005). On the salience of political issues: The problem with 'most important problem'. *Electoral Studies*, 24, 555-579. DOI:10.1016/j.electstud.2005.01.009

Table 1: Combined forecasts of BI-H and BI-M (percentages)

	2008	2004	2000	1996	1992	1988	1984	1980	1976	1972	All
Correct	91	39	89	100	100	55	100	99	99	100	87
Tie	9	43	6	0	0	45	0	1	1	0	11
Wrong	0	18	5	0	0	0	0	0	0	0	2

[*correct* if both predicted correct winner, *wrong* if both predicted wrong winner, *tie* if forecasts disagreed]

Table 2: Out-of-sample forecasts and absolute errors of BI-M (9 weeks prior to Election Day)

Election year	Actual	Predicted	Absolute error
1972	61.8	55.7	6.1
1976	48.9	49.8	0.9
1980	45.4	47.6	2.2
1984	59.2	52.8	6.3
1988	53.9	51.3	2.6
1992	46.5	47.4	0.9
1996	54.7	55.8	1.1
2000	50.3	56.4	6.2
2004	51.2	51.7	0.5
2008	46.3	49.4	3.1
		MAE	3.0

**Table 3: BI-M vs. benchmark models: Absolute errors of out-of-sample forecasts
(last three elections from 2000 to 2008, calculated by successive updating)**

Model	Approx. date of forecast	2000	2004	2008	MAE
BI-M	Late August / early September	6.8	0.8	3.1	3.6
Econometric models					
Norpoth	January	4.7	3.5	3.6	3.9
Lockerbie	May / June	10.0	6.4	4.5	7.0
Fair	Late July	0.5	6.3	2.2	3.0
Abramowitz	Late July / August	2.9	2.5	0.6	2.0
Holbrook	Late August / early September	10.0	3.3	2.0	5.1
Lewis-Beck and Tien	Late August	5.1	1.3*	3.6	3.3
Wlezien and Erikson	Late August	4.9	0.5	1.5	2.3
Campbell	Early September	2.5	2.6	6.4*	3.8
MAE (error of the typical econometric model)		5.1	3.6	2.6	3.8

* predicted the wrong winner

Table 4: Hit rate (in %) of BI-H and BI-M and the IEM vote-share and winner-take-all markets

	Election year (n=100 per election)						Mean	Mean
	1988	1992	1996	2000	2004	2008	1988-2008 (n=600)	1992-2008 (n=500)
BI-H	55	100	100	89	39	100	81	86
BI-M	100	100	100	95	82	91	95	94
IEM vote-share	67	63	100	29	85	100	74	75
IEM winner-takes-all	-	86	100	47	96	100	-	86

Table 5: BI performance and most important issues as seen by voters

Election year	% of combined BI forecasts of election winner			No. of most important issues	Problem category			No. of times voter perception changed
	Correct	Tie	Wrong		Economic	Foreign	Other	
2008	91	9	0	1	1	0	0	0
2004	39	43	18	2	1	1	0	7
2000	89	6	5	4	0	0	4	6
1996	100	0	0	4	2	0	2	3
1992	100	0	0	2	2	0	0	1
1988	55	45	0	3	2	0	1	6
1984	100	0	0	3	2	1	0	3
1980	99	1	0	2	2	0	0	2
1976	99	1	0	1	1	0	0	0
1972	100	0	0	1	0	1	0	0

Figure 1: Keyword popularity on Google Insights for Search for “Economy” and “Iraq” (2004 - 2008)

