Judgment and Decision Making, Vol. 6, No. 1, February 2011, pp. 73–88

Forecasting elections with mere recognition from small, lousy samples: A comparison of collective recognition, wisdom of crowds, and representative polls

Wolfgang Gaissmaier* Julian N. Marewski† ‡

Abstract

We investigated the extent to which the human capacity for recognition helps to forecast political elections: We compared naïve recognition-based election forecasts computed from convenience samples of citizens' recognition of party names to (i) standard polling forecasts computed from representative samples of citizens' voting intentions, and to (ii) simple—and typically very accurate—wisdom-of-crowds-forecasts computed from the same convenience samples of citizens' aggregated hunches about election results. Results from four major German elections show that mere recognition of party names forecast the parties' electoral success fairly well. Recognition-based forecasts were most competitive with the other models when forecasting the smaller parties' success and for small sample sizes. However, wisdom-of-crowds-forecasts outperformed recognition-based forecasts in most cases. It seems that wisdom-of-crowds-forecasts are able to draw on the benefits of recognition while at the same time avoiding its downsides, such as lack of discrimination among very famous parties or recognition caused by factors unrelated to electoral success. Yet it seems that a simple extension of the recognition-based forecasts—asking people what proportion of the population would recognize a party instead of whether they themselves recognize it—is also able to eliminate these downsides.

Keywords: political elections, recognition, forecasting, heuristics, wisdom of crowds.

1 Introduction

"The trouble with free elections is, you never know who is going to win", former political leader of the Soviet Union, Leonid Brezhnev, is supposed to have said once (Rees, 2006). This did not only bother Brezhnev, but also keeps polling agencies busy around the world. They usually rely on *intention-based election forecasts*, generated by interviewing large representative samples of citizens about their *voting intentions*. For instance, in Germany potential voters are typically asked which political party they will vote for in an upcoming election. The resulting responses can be used to extrapolate likely election results.

Here, we investigate how far one can get with a much simpler, almost naïve, method that does not require large and representative samples. Specifically, we test how well citizens' memories that they have heard of a party name before, that is, citizens' *mere recognition* of party names, allows forecasting the outcomes of major political elections. We compare the performance of such *recognition-based election forecasts*, computed from small and unrepresentative convenience samples of citizens, to other forecasting methods, including (i) traditional polls computed from large representative samples of citizens' voting intentions, and (ii) a simple—but typically very accurate—forecasting method that builds on the aggregated judgments of many, or the *wisdom of crowds* (Galton, 1907; Sjöberg, 2009; Surowiecki, 2004).

The article is structured as follows. First, we review previous research showing that recognition allows making accurate forecasts in many domains. Second, we explain why recognition could be an accurate predictor variable for forecasting elections and why recognition-based election forecasts could be particularly useful for forecasting smaller political parties' electoral success. Third, we introduce election forecasts based on the wisdoms of the crowds. Finally, we report and discuss a series of studies that investigate the accuracy of recognition-based election forecasts compared to forecasts based on polls of citizens' voting intentions and forecasts based on the wisdom of crowds.

Both authors contributed equally. Authorship order was determined by a coin flip. We thank Jon Baron, Mirta Galesic, Ralph Hertwig, Konstantinos Katsikopoulos, and Henrik Olsson for many helpful comments.

*Max Planck Institute for Human Development, Harding Center for Risk Literacy, Lentzeallee 94, 14195 Berlin, Germany. Email: gaissmaier@mpib-berlin.mpg.de

†Max Planck Institute for Human Development, Center for Adaptive Behavior and Cognition, Lentzeallee 94, 14195 Berlin, Germany. Email: marewski@mpib-berlin.mpg.de.

‡IESE Business School, Barcelona, Spain

Figure 1: Goodman and Kruskal's (1954) gamma computed between the frequency of mentions of 25 parties in the newspaper "Tagesspiegel" in the period between 5 months 16 days and 16 days prior to the National elections 2005, the number of votes won by 25 parties in that election, and the number of participants who recognized the name of a party 16 days prior to the election (cor: correlation). These correlations show that the unknown criterion (here: the election result) is reflected by a mediator (here: the newspaper "Tagesspiegel"). The mediator makes it more likely for a person to encounter alternatives with larger criterion values than those with smaller ones (e.g., the press mentions more successful political parties more frequently). As a result, the person will be more likely to recognize alternatives with larger criterion values than those with smaller ones, and, ultimately, recognition judgments can be relied upon to infer the criterion (here: the success of parties in political elections).

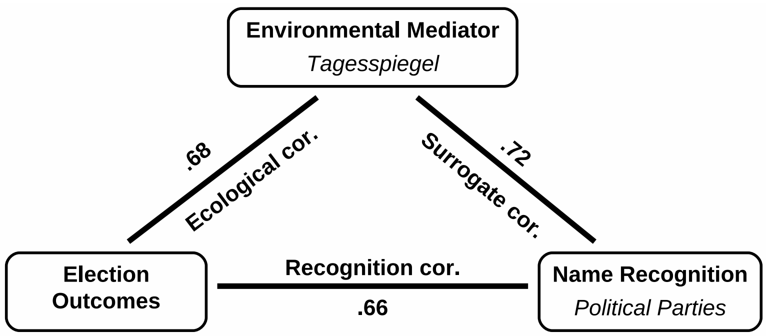

1.1 The predictive power of recognition in forecasting

Why would recognition be useful for forecasting in general? A major reason is an ecological one (Goldstein & Gigerenzer, 2002; Hertwig, Herzog, Schooler, & Reimer, 2008; Schooler & Hertwig, 2005): The press, the internet, and other *environmental mediators* make it likely that we will encounter *objects* (e.g., tennis players, cities, universities) that score high on a *criterion* of interest (e.g., success in sports, size of cities, quality of universities) more frequently than those that score low. As a result, objects with high criterion values are more likely to be recognized. Thus, when making forecasts, we can rely on recognition to predict which objects are likely to score high on the criterion.

The simple forecasting strategy to bet that objects that are recognized by more people will score higher on a criterion of interest is also known as the *collective recognition heuristic* (e.g., Borges, Goldstein, Ortmann, & Gigerenzer, 1999; Herzog & Hertwig, 2011): Count how many people recognize each of N objects, and infer the n recognized objects to score a larger value on the criterion than the $N - n$ unrecognized ones. It has been shown that people's collective recognition allows for making accurate forecasts in many domains. The outcomes of Wimbledon tennis matches, for instance, can be predicted by simply betting that those players who are recognized by most people will win (Scheibehenne & Bröder, 2007; Serwe & Frings, 2006). Such naïve recognition-based forecasts were more accurate than Association of Tennis Professionals rankings or Wimbledon seeds. Other domains where recognition makes good predictions include forecasts about the sizes of cities (Goldstein & Gigerenzer, 2002; Reimer & Katsikopoulos, 2004), the quality of universities (Hertwig & Todd, 2003), the fortunes of billionaires (Hertwig et al., 2008), and the success of soccer teams in championships (Pachur & Biele, 2007).

We have good reasons to believe that collective recognition will also allow forecasting elections. For one of the elections that we studied (German National Elections 2005, see below), Figure 1 shows that there are substantial correlations between (i) election results, (ii) the frequency of newspaper mentions, and (iii) the number of people who recognized a party's name. Thus, before we test collective recognition in more detail against other models, this already is a first illustration that the domain of elections is principally suited for collective recognition (see also Marewski, Gaissmaier, Schooler, Goldstein, & Gigerenzer, 2010). In the next section, we argue in more detail why we believe that recognition could allow making accurate election forecasts.

1.2 Four reasons why recognition may help to forecast elections

1.2.1 Robustness of recognition with respect to the characteristics of the citizens in the sample

First, we suspect that recognition-based election forecasts are relatively robust to the characteristics of the sample used to compute the forecasts. For instance, Scheibehenne and Bröder (2007) found that both experts and laypeople's recognition of tennis players' names

https://doi.org/10.1017/S1930297500002102 Published online by Cambridge University Press

yielded almost equally good predictions of the outcomes of Wimbeldon tennis matches, although laypeople knew only very little about tennis and recognized, on average, only about one fifth of the names that the experts recognized. Likewise, when it comes to deriving election forecasts, one may expect samples of people's recognition of party names to be more robust to sampling biases than samples of people's voting intentions. To illustrate this, in a sample of German psychology students, the proportion of voters for left-wing parties will be overrepresented. Hence, election forecasts computed from these students' voting intentions will be biased towards the left-wing parties. German psychology students, however, are exposed to largely the same environmental mediators (e.g., TV, radio, newspaper, Internet) as the rest of the electorate. As a result, these students' recognition of party names is likely to be more representative of the electorate than the same students' voting intentions.[1]

1.2.2 Robustness of recognition with respect to the influence of psychological variables

Second, even though a sense of name recognition can be easily induced (e.g., by advertising firms or politicians placing election ads in an election), once a name *is* recognized, the recognition of this name is comparatively robust against the influence of other psychological variables. For instance, a sense of recognition is remarkably lasting and does not decline as much with age as recall memory (e.g., Craik & McDowd, 1987). At the same time, recognition is easily accessible, and likely to emerge on the mental stage earlier than other information a person may recall about a name (e.g., Pachur & Hertwig, 2006). Shepard (1967) tried to quantify the human capacity of recognition memory. In his experiment, subjects were shown 612 pairs of photographs. In a paired comparison task with new pictures, subjects' recognition accuracy was as high as 99%. Even when Standing (1973) increased the number of pictures to 10,000, subjects were able to tell with a very large accuracy which pictures they had seen before and which not. Voting intentions, in contrast, can be influenced by a host of other psychological variables, such as a person's momentary political preferences or her mood. In fact, in many democracies some proportion of swing voters end up voting differently then they declare in election surveys conducted beforehand. Such changes in voting intentions can systematically bias the accuracy of intention-based election forecasts, but should affect to a lesser extent the accuracy of recognition-based forecasts, as voters may be able to easily change their intentions on a day-by-day basis, but are unlikely to erase a sense of recognition from their minds.

1.2.3 Robustness of recognition with respect to sample size in forecasts for smaller parties

Third, in order to be accurate, recognition-based forecasts are likely to require smaller sample sizes of interviewed citizens than intention-based forecasts. For instance, in Germany, there are often between 1 or 2 dozen parties competing in elections. Yet the vast majority of votes, typically between 90 and 95%, will go to the 4 or 5 larger German parties, with only few votes being casted for the remaining smaller parties. Correspondingly, in surveys of voting intentions very few people (if any at all) will declare that they intend to vote for one of the smaller parties, resulting in very few observations that could be used to compute intention-based election forecasts for these smaller parties. As a result, intention-based forecasts for these smaller parties require very large samples of interviewees in order to be accurate, making such forecasts costly. This is, perhaps, also one reason why pollsters usually refrain from publishing polls for such small parties. In contrast, when interviewing Germans about their recognition of these smaller parties, many will still recognize their names, which could allow making accurate forecasts about small parties' electoral success even when the sample of interviewed voters is small. Put differently, when it comes to forecasting smaller parties electoral success, recognition-based forecasts may be more robust with respect to the sample size than intention-based ones.

1.2.4 The role of recognition in decision making and voting

Fourth, recognition plays an important role in decision making (for a recent review, see Pachur, Todd, Gigerenzer, Schooler, & Goldstein, in press): To illustrate this, a sense of recognition can determine what people like (e.g., Zajonc, 1968), which consumer products they prefer (e.g., Coates, Butler, & Berry, 2004, 2006), or which companies and cities they believe to be big (Goldstein & Gigerenzer, 2002; Goldstein, 2007; Hertwig et al., 2008; Hilbig, 2008; Hilbig & Pohl, 2008; Marewski, Gaissmaier, Schooler, et al., 2010; Newell & Fernandez, 2006; Pachur, Bröder, Marewski, 2008; Pohl, 2006; Volz et al., 2006). And in the political science and polling literatures it has long been known that recognition plays an important role in voting. For instance, there is evidence that recognition influences candidate preference (e.g., Goldenberg & Traugott, 1980). In fact, recognition could actually help voters to cast their ballots in a smart way even when they know little about the candidates and parties competing in an election. Voters rely

[1] We thank Ralph Hertwig for pointing out why recognition may be less prone to sampling biases than voting intentions for forecasting elections.

https://doi.org/10.1017/S1930297500002102 Published online by Cambridge University Press

on simple rules of thumb, or *heuristics*, to make decisions (Gigerenzer, 1982, 2007; Jackman & Sniderman, 2002; Kelley & Mirer, 1974; Sniderman, 2000; Todorov, Mandisodza, Goren & Hall, 2005; Wang, 2008; see also Popkin's 1994). In deciding how to vote, especially voters who know little about political issues could go with the heuristic to choose recognized candidates and parties. After all, voters do not only take the desirability of candidates or parties into account, but also their likelihood of being elected (Stone & Abramowitz, 1983), and using this heuristic could help even ignorant voters to identify likely winners or, at least, to eliminate losers from consideration (see Marewski, Gaissmaier, Schooler et al., 2009, 2010, for corresponding evidence[2]). In Germany and many other countries, candidates and parties receive funding as a function of their past electoral success, which in turn may influence both their name recognition and their success in future elections. And for the United States, the political science literature documented that the advantages of incumbency, including better campaign financing, greater name recognition, and more positive voter evaluations, are critical factors affecting voting decisions (e.g., Abramowitz, 1975; Campbell, Alford, & Henry, 1984; Goldenberg & Traugott, 1980; Jacobson, 1987; Mann & Wolfinger, 1980; Miller & Krosnick, 1998). This literature thus suggests that name recognition may allow forecasting elections.

1.3 Wisdom-of-crowds-forecasts: Another simple forecasting method

Besides recognition, there are other techniques that allow forecasting elections in a simpler way than traditional polls of voting intentions. One such forecasting technique is based on the wisdom of crowds, which was investigated more than 100 years ago by Sir Francis Galton, who visited a livestock fair where villagers estimated the weight of an ox. Galton was surprised to find that their median and mean average estimates were only 9 and 1 pounds, respectively, off the actual weight of 1198 pounds (Galton, 1907). Subsequently, it was repeatedly shown for many domains that averaging the predictions of many can improve the overall performance of forecasts about future events or unknown quantities (e.g., Armstrong, 2001; Clemen, 1989; Hogarth, 1978; Johnson, Budescu, & Wallsten, 2001; Surowiecki, 2004; Timmermann, 2006; Wolfers & Zitzewitz, 2004).

In elections, Sjöberg (2009) showed that the wisdom of crowds actually allowed for more successful forecasts than polls, making it a strong competitor to recognition. Another reason why such *wisdom-of-crowds-forecasts* may represent a strong competitor to recognition is that wisdom-of-crowds-forecasts of elections may actually be partially based on recognition, combing recognition with other useful information. To generate wisdom-of-crowds-forecasts, one asks citizens to guess the election result; for instance, by rank ordering parties according to the number of votes a citizen believes the parties will win. These individual hunches are averaged across citizens, and the average is used as a prediction of the election outcome. In past studies, we (Marewski, Gaissmaier, Schooler, et al., 2010) have provided evidence that citizens rely heavily on their recognition of party names to generate such hunches about election outcomes, betting that the parties they recognize will win more votes than those they do not. In comparisons of recognized parties, in turn, citizens tend to rely on other information they may recall about the parties, such as the parties' political agenda, publically available polls, or the parties' past electoral success. To the extent that this other information reflects the likely election result, wisdom-of-crowds-forecasts that take this information into account may turn out to be more accurate than forecasts that rely on collective recognition alone.

For instance, based on publically available polling information, citizens may be able to accurately forecast the rank order of votes for the 4–5 larger German parties, using their recognition of party names to forecast the rank order of votes for the remaining smaller parties. For these remaining parties, forecasts based on collective recognition will thus generate similar rank orders of predicted votes as wisdom-of-crowds-forecasts; however in contrast to the wisdom-of-crowds-forecasts, the recognition-based forecasts are unlikely to reflect the rank order of votes the 4–5 largest German parties will win, because most Germans will recognize the names of all of these parties.

Moreover, while wisdom-of-crowds forecasts and recognition-based forecasts are likely to be similar for smaller political parties, they do not need to be identical: Also for forecasts about the smaller parties, wisdom-of-crowds-forecasts may enjoy an advantage over recognition. In many democracies, there are a couple of smaller parties that are highly recognized although only few people will vote for them, as is often the case for radical right-wing parties. Recognition-based forecasts may thus forecast unrealistically large numbers of votes for these small, highly-recognizable parties.

[2] Marewski, Gaissmaier, Schooler, et al. (2009, 2010) provided evidence to suggest that voters rely on their recognition of political parties' names to forecast the outcomes of German political elections. However, Marewski, Gaissmaier, Schooler, et al. focused on how individuals make election forecasts, and not on the collective recognition heuristic, or on forecasting techniques in general. As such, they also did not evaluate *how well* recognition predicts election outcomes by comparing recognition-based election forecasts against other forecasting models.

2 Study methods

2.1 Overview of the studies

To test how well recognition allows forecasting elections in comparison to standard polls and the wisdom-of-crowds-principle, we studied four important elections in Germany, which is the largest democracy in the European Uninon[3]: The 2004 parliamentary elections in the federal state of Brandenburg, the 2005 parliamentary elections in the federal state of North Rhine-Westphalia, and the 2005 and 2009 German national elections. For the first three elections, we reanalyzed recognition data that had originally been collected by Marewski, Gaissmaier, Schooler, et al. (2009; 2010). For the fourth election, we ran a new study. This new data allowed us to run additional analyses that were not possible in the reanalyses.

Participants in all studies were small convenience samples of university students or pedestrians interviewed on the streets—samples most professional pollster would deem lousy. In all studies, in a *recognition task*, participants from these samples were either given lists of parties' names in a questionnaire (Studies 1, 2, and 4) or presented parties' names on a computer screen (Study 3). The names were always randomly ordered. For each name, people were asked whether they had heard of or seen it before participating in the study. Participants could answer with yes or no. We will refer to these binary decisions as *recognition judgments*. In Studies 1 to 3, in a *voting intention task* participants were asked for which party they intend to vote in the upcoming election, using the question format that is regularly employed by German polling institutions.[4] Participants answered by writing down the party name or its abbreviation.[5] We will refer to these responses as *observed voting intentions.* Completing these tasks took only a few minutes.

All studies also included a *prediction task*, which we used to construct wisdom-of-crowds-forecasts. In this task, people were asked to forecast which party would receive more votes. To this end, participants were either asked to rank all parties according to their prediction of the election outcome (Studies 1, 2, and 4) or to predict for all possible comparisons of two parties which one would win (Study 3). The order of parties and the order of pairs of parties were randomized.

Study 4 aimed at replicating the results of our reanalyses of Studies 1 to 3, but it also had two important extensions. First, the voting intention task typically used by polling institutions and employed by us in Studies 1 to 3 yields only one observation per interviewee, that is, one voting intention for *one* party, given by one subject. In contrast, our recognition task entails gathering several observations per interviewee, namely one recognition judgment *for each of the N* parties competing in an election, given by one subject. To rule out that the possibility that this difference in the number of observations is responsible for potential differences between the accuracy of intention-based election forecasts and recognition-based ones, we extended the voting intention task in Study 4. Rather than eliciting solely a single voting intention, we additionally asked participants to rank order the remaining parties according to their voting preferences. Specifically, we asked participants to rank the party they intended to vote for at position one. All other parties were to be assigned a lower rank in the order of their preferences. This *extended voting intention task* yields one observation per party, and as such, an equal amount of observations as the recognition judgment task. We will refer to these rankings as *observed voting intention rankings.* Besides comparing recognition-based forecasts to intention-based ones, the extended voting task allows us to additionally assess how well intention-based forecasts computed from aggregating intention rankings predict elections compared to intention-based forecasts computed from eliciting just one voting intention (i.e., the party ranked above all others).

As a second extension of Study 4, we tried to push the recognition principle a little further. As mentioned above, for the 4–5 larger German parties and other highly recognizable parties (e.g., certain extreme left-wing or right-wing parties), recognition-based forecasts face the problem that these parties are recognized by everyone, making it difficult to predict which of these parties will win an election. In this case, recognition is said to not *discriminate* between the parties. To counter this discrimination problem, in a *recognition estimation task* we asked participants to estimate how many out of 100 people would recognize each party. We hoped that these *subjective recognition estimates* would exhibit a larger variance than

[3]Like most other European democracies, Germany is a multi-party system, in which approximately 15 to 30 parties compete on both the national and the federal level. In most German states as well as on the national level, every 4 years, each citizen has two votes; one for a direct candidate who will represent the person's voting district and a second for a party, representing a list of candidates. Direct candidates are typically affiliated with one of the parties and are elected into Parliament if they win the most votes in their voting district. If a party is elected into Parliament, then, depending on its proportion of votes, a number of the candidates from its list enter Parliament.

[4]The precise phrasing of the voting intention question was: "The election takes place on Sunday the X^{th}, for which party will you vote?" In all studies, the phrasing included the name and date of the election. The precise phrasing of the recognition question was: "Do you recognize this party name, that is, have you heard or seen it before participating in this study?"

[5]There were a few participants who either indicated not to vote or to cast an invalid ballot, or who simply left the answer to this question blank.

https://doi.org/10.1017/S1930297500002102 Published online by Cambridge University Press

recognition judgments alone, which in turn, may allow for better discriminating between such parties.

2.1.1 Study 1: State elections in Brandenburg 2004

At two dates, 14 days and 1 day before the election, we invited pedestrians in the downtown areas of the Brandenburgian cities of Potsdam and Werder to fill out a questionnaire. The only criterion to select participants was that they were eligible to vote. Of 246 recruited participants, 172 completed the questionnaire (70%; 55% female; mean age 38 years, *SD* = 14.7). All participants were at least 18 years old (voting age in Germany). They were paid €5 ($7).

2.1.2 Study 2: State elections in North Rhine-Westphalia 2005

Fifty-nine university students from Berlin, Germany, (43% female, mean age 26 years, *SD* = 3.6) filled out a questionnaire 3 to 11 days before the election. About half of them completed the questionnaire in our lab and received €5 ($7) for their participation; the other half worked on it in a university class. All participants had to be at least 18 years of age, but were unlikely to be eligible to vote in North Rhine-Westphalia as they lived about 400 km away from that state.

2.1.3 Study 3: German national elections in 2005

Sixty-six residents of Berlin, Germany, most of them students (52% female; mean age 26 years, *SD* = 3.7), participated in the study. They were recruited from the subject pool of our research institution. All participants were at least 18 years old and eligible to vote. They were paid €25 ($37). The assessment took place 16 days prior to the election and was part of larger study.

2.1.4 Study 4: German national elections in 2009

Thirty-four residents of Berlin, Germany, most of them students (56% female; mean age 25 years, *SD* = 3.0), completed a computerized survey in our laboratory during the week before the election. They were recruited from the subject pool of our research institution. All participants were at least 18 years old and eligible to vote and participated as part of other studies without being paid extra for it. In addition to the tasks employed in the other studies, they completed a recognition estimation task, in which they had to estimate how many out of a 100 randomly drawn people would recognize a party, as well as an extended voting intention task, in which they had to rank all parties in order of their preferences, assigning the top rank to the party they actually intended to vote for. The order of all tasks was randomized.

2.2 Forecasting Models

To test how good recognition does in forecasting elections, we tested a total of three classes of models: Recognition-based forecasts, intention-based forecasts, and wisdom-of-crowds-forecasts.

2.2.1 Recognition-based forecasts

Prior to each election we counted how many participants recognized each party's name and used this count to predict the rank order of the number of votes the parties would win (REC/*basic*). This recognition-based forecasting model corresponds to the collective recognition heuristic used in earlier studies for predicting sport events and the performance of stocks (e.g., Borges et al., 1999; Serwe & Frings, 2006; Herzog & Hertwig, 2011). In Study 4, we additionally tested recognition-based forecasts generated from participants' subjective estimates how many out of 100 randomly drawn people would recognize each party. We averaged these subjective recognition estimates across participants and used this average to forecast the rank order of the number of votes the parties would win in the election (REC/*extended*).

2.2.2 Intention-based forecasts

To evaluate the performance of naïve recognition-based forecasting models, we constructed benchmark models that simulated the representative sampling of voting intentions. As upper benchmark, we simulated intention-based forecasts with samples of size 20 to 1,000 in steps of 20 drawn from the *actual* election results. For each sample size, we repeated this procedure 10,000 times. That is, we generated *perfectly* representative samples of how voters actually decided (INT/*representative*). However, real intention-based forecasts can suffer from both sampling error and swing voters who vote differently from what they declare in surveys. To make our intention-based forecasts more realistic, we ran additional simulations where we randomly reassigned 5% of voters of each of the parties to have voted for a different party—as if they had reconsidered their choice. These simulations were also repeated 10,000 times for sample sizes 20 to 1,000 in steps of 20 (INT/*representative + swing voters*).

As a lower benchmark, we also computed intention-based forecasts from our study participants' observed voting intentions (INT/*study sample*). This model INT/*study sample* not only enabled us to compare the performance of intention-based forecasts computed from lousy samples to the performance of recognition-based forecasts computed from the same lousy samples, but also allowed us to assess how little representative our sample of participants' voting intentions was of the German electorate's votes.

Finally, for Study 4, we additionally computed intention-based forecasts from participants' observed voting intention rankings. To do so, we averaged these rankings across participants and used this average to forecast the rank order of the number votes the parties would win (INT/*study sample rankings*).

2.2.3 Wisdom-of-crowds-forecasts

Based on the prediction tasks in which we had asked people to predict which parties would gain more votes than others, we constructed wisdom-of-crowds-forecasts. Specifically, we averaged the predicted ranks of electoral success across study participants in each of the study and used these averages to forecast the election outcomes (WIS).

2.3 Performance Measures

2.3.1 Ordinal predictions

Just as the collective recognition heuristic, also all other simple forecasting models considered here make ordinal predictions of election outcomes (i.e., REC/*basic;* REC/*extended;* WIS). We therefore compared all models' ability to predict the rank order of votes the political parties received. To do so, we generated all pairwise comparisons between all parties. For REC/*basic*, across all pairs we counted how often the party that won more votes in the election was the one that was recognized by more people. Likewise, for REC/*extended*, across all pairs we counted how often the party that won more votes in the election was the one that the participants of Study 4 had estimated to be, on average, recognized by more people. For the four intention-based models, we counted how often the party that won more votes was the one that had received more voting intentions, using the simulated voting intentions (INT/*representative* and INT/*representative + swing votes*), the observed voting intentions (INT/*study sample*), and the averaged observed voting intention rankings (INT/*study sample rankings*), respectively. For the WIS model, we counted how often the party that won more votes was the party that was assigned the better rank, averaged across participants. Whenever there was a tie, either because both parties were recognized by the same number of people or because there were equally many voting intentions for both parties or because the mean predicted rank was identical, the models made random guesses. The *accuracy of the forecasts* is the resulting proportion of correct predictions, computed across all comparisons between two parties.[6]

2.3.2 Predictions of shares of votes

Typically, the goal of election forecasts is not only to predict an ordinal rank order but also to forecast shares of votes. The predictor variables used in the simple forecasting models evaluated here (i.e., REC/*basic;* REC/*extended;* WIS) could, in principle, be incorporated in corresponding *estimation models*, for instance, by assigning weights to them that translate ordinal ranks into shares of votes. It is beyond the scope of this paper to systematically evaluate which of many plausible estimation models (e.g., including different weights and functional forms) is most accurate; however, we will also present a smaller subset of additional analyses that allow exploring how well recognition as a predictor variables could, at least in principle, allow for forecasting shares of votes. In doing so, we will focus on the shares of votes the smaller political parties gain: As explained above, it takes very large samples to predict shares of votes for these smaller parties based on surveys of voting intentions, such that a simpler alternative forecasting technique may actually help here. Recognition, in contrast, may allow generating accurate forecasts based on small samples, and could thus be particularly useful when forecasting the small parties' success. Much the same can be said with respect to simple forecasts based on the wisdom of crowds: As we have explained above, these forecasts are likely to be partially based on recognition; correspondingly, also they may help forecasting the smaller parties' electoral success.

2.3.3 Large versus small parties

As recognition may be particularly useful for forecasting smaller parties electoral success, all ordinal forecasts were computed separately for both the complete set of all parties and for a subset of small parties. Smaller parties were those that were not represented in the German national Parliament at the time of the election.[7] (To enter the national Parliament, a party needs to gain more than 5% of the votes in the national elections.) There were 15 parties competing in Brandenburg, 24 in North Rhine-Westphalia, 25 in the national elections 2005, and 27 in the national elections 2009. The subset of small parties consisted of 10, 19, 19, and 21 parties, respectively.

[6] Other accuracy measures (Kendall and Spearman rank correlations) yielded the same patterns of results.

[7] By our definition, the large parties were CDU/CSU, DieLinke/PDS, FDP, GRÜNE, and SPD, all other parties were considered small. We also ran the analyses using other criteria to define the subset of smaller parties. The pattern of results remained the same.

Figure 2: Forecasting German elections with mere recognition with seven different forecasting models: (REC/*basic*) recognition-based forecasts using averaged recognition judgments from our study participants; (INT/*representative*) intention-based forecasts using a simulated, perfectly representative sample of voters (means + *SD*); (INT/*representative* + *swing voters*) intention-based forecasts using a simulated, perfectly representative sample of voters, but letting 5% of voters of each party reconsider their choice by randomly reassigning them to have voted for a different party (means + *SD*); (INT/*study sample*) intention-based forecasts computed from the observed voting intentions of our study participants; (WIS). Forecasts based on the mean predicted ranks by our study participants. Two forecasting models could only be computed for Study 4: (REC/*extended*) recognition-based forecasts based on participants' subjective estimates how many out of 100 randomly drawn people would recognize each party, averaged across participants; (INT/*study sample rankings*) intention-based forecasts based on average observed voting intention rankings provided by the participants for each of the parties. All results are depicted separately for the subset of small parties, which are not represented in German Parliament and for which usually no polls exist (upper panels), and for all parties (lower panels). Note that a proportion correct of 0.5 represents chance level, that is, the accuracy that would be achieved by randomly guessing in all paired comparisons between two parties. Further note that in panel IIa, REC/*basic* and WIS are based on the same sample size and are just moved apart for reasons of readability, and the same is true for REC/*extended* and WIS in panels IVa and IVb.

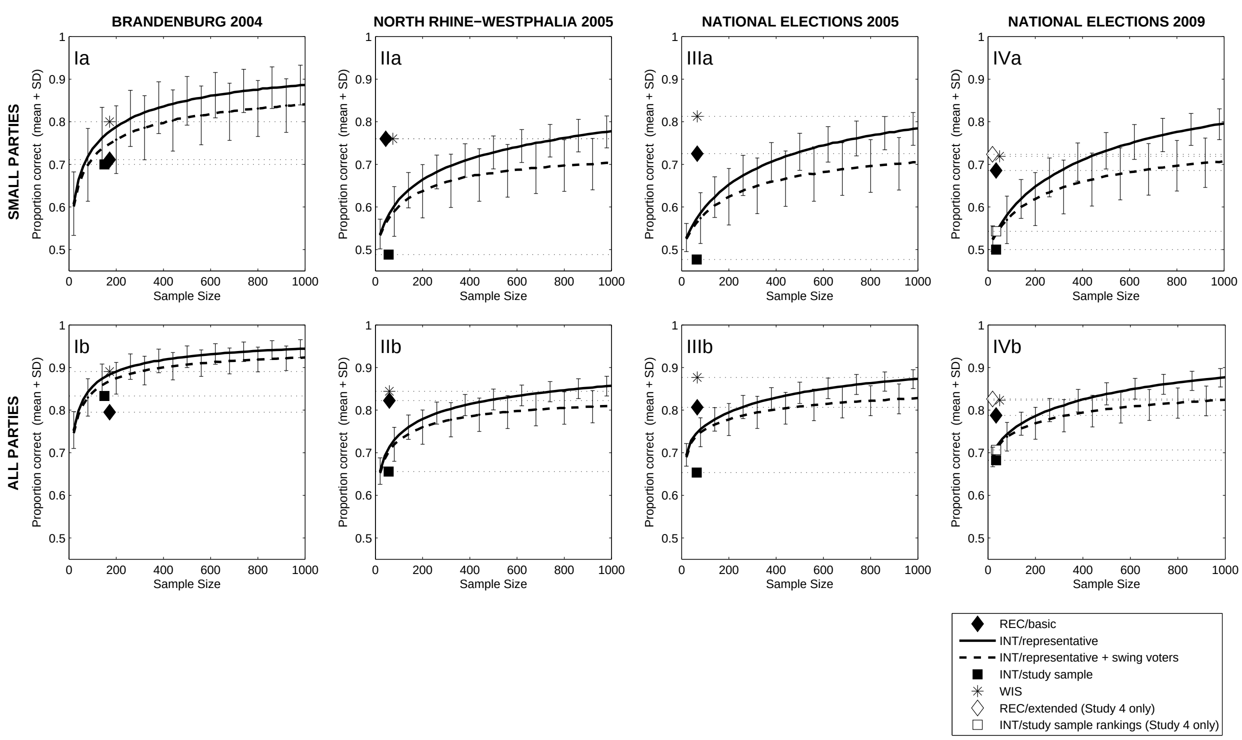

https://doi.org/10.1017/S1930297500002102 Published online by Cambridge University Press

3 Results and discussion

3.1 Ordinal predictions

Figure 2 shows the proportion of correct recognition-based forecasts, intention-based forecasts and forecasts based on the wisdom of crowds. First, intention-based forecasts computed from the convenience samples (INT/*study sample*) were the least accurate, illustrating that the study samples were indeed unrepresentative of how German voters decided in the election (with the exception of Brandenburg, which we will discuss separately below). Just to give one example of how different the electoral preferences of our samples were in comparison to the general population, consider Study 4: Here, 44.1% of participants would have voted for the Green party, while this party only received 10.7% of the votes in the general population. Importantly, as comparing REC/*basic* and INT/*study sample* shows, recognition-based forecasts, computed from the very same unrepresentative samples, tended to fare considerably better than the intention-based ones, suggesting that recognition is indeed a predictor variable that is fairly robust to the characteristics of the citizens included in the sample.

Importantly, this difference between intention-based and recognition-based forecasts from the convenience samples does not stem from a difference in number of observations. Recall that in Study 4 we had additionally asked participants to rank all parties according to their voting preferences (INT/*study sample rankings*). Although these complete voting intention rankings notably improved intention-based forecasts based on the convenience samples, these forecasts are still much inferior compared to recognition-based forecasts from the same unrepresentative samples (panels IVa and IVb).

Second, as comparisons of REC/*basic*, INT/*representative* and INT/*representative + swing voters* reveal, unrepresentative recognition-based forecasts can compete with intention-based forecasts computed from *perfectly representative* samples, especially for the subset of smaller parties (see upper panels). One reason for this is that few people vote for the small parties, which makes it necessary to survey extremely large samples to get reliable estimates for intention-based forecasts. For instance, as Figure 2 shows, interviewing about 1,000 individuals is still not enough to generate accurate election forecasts for small parties based on perfectly representative samples. In comparison, recognition does relatively well, even when based on very small, unrepresentative samples. In short, when it comes to forecasting the smaller parties' electoral success, recognition-based forecasts seem to be more robust with respect to the sample size than intention-based ones.

In fact, as comparisons of REC/*basic*, INT/*representative* and INT/*representative + swing voters* in the set of all parties show (see lower panels), unrepresentative recognition-based forecasts were generally most likely to reach the level of accuracy of perfectly representative intention-based forecasts when the sample size of surveyed individuals was small. For instance, in panel IIb (North Rhine-Westphalia, all parties), the mean accuracy attainable with mere name recognition exceeded the mean accuracy of representative intention-based election forecasts until up to a sample sizes of about 400 surveyed voters.

Perhaps most interestingly, also for relatively large sample sizes (e.g., 1000 in Panel IIa, 700 in Panel IIIa, and 500 in Panel IVa), the mean accuracy of unrepresentative recognition-based election forecasts fell within the range of 1 standard deviation of the accuracy of perfectly representative intention-based election forecasts (with the notable exception of panels Ia and Ib, Brandenburg). Note that this relative advantage of recognition-based election forecasts emerged even when participants knew very little about the election, as is the case in panels IIa and IIb, where all study participants lived in a different federal state than the one in which the election took place (North Rhine-Westphalia).[8]

Third, WIS outperformed REC/*basic* in almost all cases, most likely because people are able to rely on other information beyond mere recognition when ranking two or more parties they recognize, which REC/*basic* cannot do. Interestingly, forecasts based on participants' averaged estimates how many out of 100 randomly drawn people would recognize each party (REC/*extended*) were basically indistinguishable from WIS. The improvement observed from REC/*basic* to REC/*extended* from the same convenience sample (panels IVa, IVb) suggests that people seem to be able to successfully discriminate between highly recognizable parties (e.g., large parties, radical parties) when estimating population recognition rates, and that it is this additional discrimination that is responsible for this increment in performance.

Finally, REC/*basic* was not competitive in comparison to intention-based forecasts in Brandenburg. We do not know why this result emerged; a plausible explanation for it may be that in Brandenburg only 15 parties competed against each other, as opposed to 24, 25, and 27 parties in the other three elections. This comparatively small number of competing parties may have boosted the accuracy of intention-based forecasts, as people's votes—and hence their voting intentions—are divided among fewer parties, making intention-based forecasts more robust to variation in the size and compo-

[8] Note that the large parties competing in German elections tend to be largely the same in different German states. However, the smaller parties vary more strongly across states.

sition of the sample of voters being drawn. In fact, as can be seen in Figure 2, it is not so much the accuracy of REC/*basic* that differed across the elections, but more the accuracy of the intention-based forecasts that was particularly high in Brandenburg. In particular, REC/*basic* achieved an accuracy of 0.80 in Brandenburg (all parties), which is basically identical to its accuracy in the other elections ranging from .79 (National Elections 2009) to .82 (North Rhine-Westphalia 2005). To compare, INT/*representative* with a sample size of 1,000 achieved an accuracy of .94 in Brandenburg (all parties), which is substantially above its accuracy in the other elections ranging from .86 (North Rhine-Westphalia 2005) to .88 (National Elections 2009).

If our explanation for the relative boost in performance of intention-based forecasts in Brandenburg is correct, then this suggests that the usefulness of REC/*basic* may be limited to elections where many parties are competing against each other. (Unfortunately, we did not test REC/*extended* in Brandenburg, so that we do not know whether the same conclusion applies to this second recognition-based forecasting model, which, as Figure 2 shows, turned out to be quite accurate, both in comparison to REC/*basic* and the intention-based forecasts in the 2009 German national elections.)

3.2 Predictions of shares of votes

To explore the continuous relation between election results on the one hand and the forecasts made by the different models on the other, we log-transformed the election results and the sampled voting intentions (Figure 3). (The log-transformation helps to visualize the data for the very small parties.) The three rows show three different model classes: Panels A show the predictions of REC/*basic* based on the convenience samples; for the German national elections 2009, panel A additionally shows predictions of REC/*extended*. Panels B show the predictions of the most accurate intention-based model, INT/*representative*, based on sample sizes of 1,000. As the predictions of INT/*representative* vary as a function of the voting intentions included in the sample being drawn in our simulations, we show 4 random draws of 1,000 voting intentions for INT/*representative*, this way illustrating the variation observed between different draws. Finally, panels C show the predictions for WIS, based on the same convenience samples as REC/*basic*. (Note that the x-axis is reversed in panels C: smaller numbers indicate more successful ranks.)

Panels A illustrate that REC/*basic* does basically not discriminate among larger parties, as all of them are recognized by about 100% of our participants. Sampling intentions, on the other hand, works better the larger the party (panels B). More precisely, sampling intentions of 1,000 individuals drawn from a representative population works pretty well until the share of votes of a party is smaller than about 1%, which is when the correlation between sampled intentions and election outcomes starts to break down. Additionally, in all elections except for Brandenburg 2004, sampling voting intentions bears a substantial risk of not at all observing voting intentions for particular parties. In Brandenburg, in contrast, voting intentions are most often observed for all parties in the race, even for the smallest ones. The reason for Brandenburg 2004 being an exception is likely to be the same we discussed above: There were fewer parties competing in the Brandenburg election than in the other elections (i.e., 15 parties in Brandenburg vs. 24 to 27 parties in the other elections), resulting in people's votes—and hence their voting intentions—being divided among fewer parties, which increases the chance to observe a voting intention for any particular party.

Comparing the scatter plots for REC/*basic* (panels A) with wisdom-of-crowds-forecasts (panels C; WIS) from the very same convenience samples reveals that wisdom-of-crowds-forecasts are generally better able to differentiate between parties. This holds true not only for the large parties but also for the small parties, although to a lesser degree. Put differently, the predictions of REC/*basic* and WIS are indeed more similar for the small parties than for the large parties; yet, WIS still provides a better reflection of the distribution of votes than REC/*basic* even for the small parties. However, as panel A shows for the national elections 2009, REC/*extended* can differentiate between parties as well as WIS. It can be nicely seen that REC/*extended* is able to eliminate the downsides of REC/*basic*, for instance by correcting unrealistically high forecasts for parties that are small, yet recognized by many people for reasons unrelated to electoral success (such as radicalism).

4 General discussion

Much research centers on forecasting the outcomes of political elections (see e.g., Campbell & Lewis-Beck, 2008; Lewis-Beck & Rice, 1992, Sigelman, Batchelor, & Stekler, 1999, for overviews). We investigated whether peoples' mere recognition of party names helps forecasting the results of political elections. As we have shown for major German elections, at least for smaller political parties recognition-based election forecasts (i.e., REC/*basic*; REC/*extended*) can be as accurate as interviewing voters about their voting intentions. In contrast to surveys of voting intentions, recognition-based election forecasts seem to be less in need of large representative samples of voters in order to be reasonably accurate. Rather, they can be computed from small, lousy samples, illustrating

Figure 3: Visual inspection of the continuous relation between election results in all four elections on the one hand and recognition-based forecasts (panels A: REC/*basic*, & REC/*extended* for national elections 2009), intention-based forecasts (panels B: INT/*representative*), and wisdom-of-crowds-forecasts (panels C: WIS), respectively, on the other. The scatter plots showing the intention-based forecasts (INT/*representative*) represent four random draws with N = 1,000 each. The scatter plots showing the recognition-based (REC/*basic*, & REC/*extended* for national elections 2009) and the wisdom-of-crowds-forecasts (WIS) represent the actual study samples with varying sample sizes, indicated on the X-axes labels. Note that both the election results and the sampled voting intentions are depicted using a logarithmic scale. For wisdom of crowds (WIS), the X-axis is reversed, as lower ranks indicate more success. The dashed horizontal lines roughly represent the split between large and small parties that we have applied, as it represents the 5% threshold that is required to enter both national and federal parliaments.

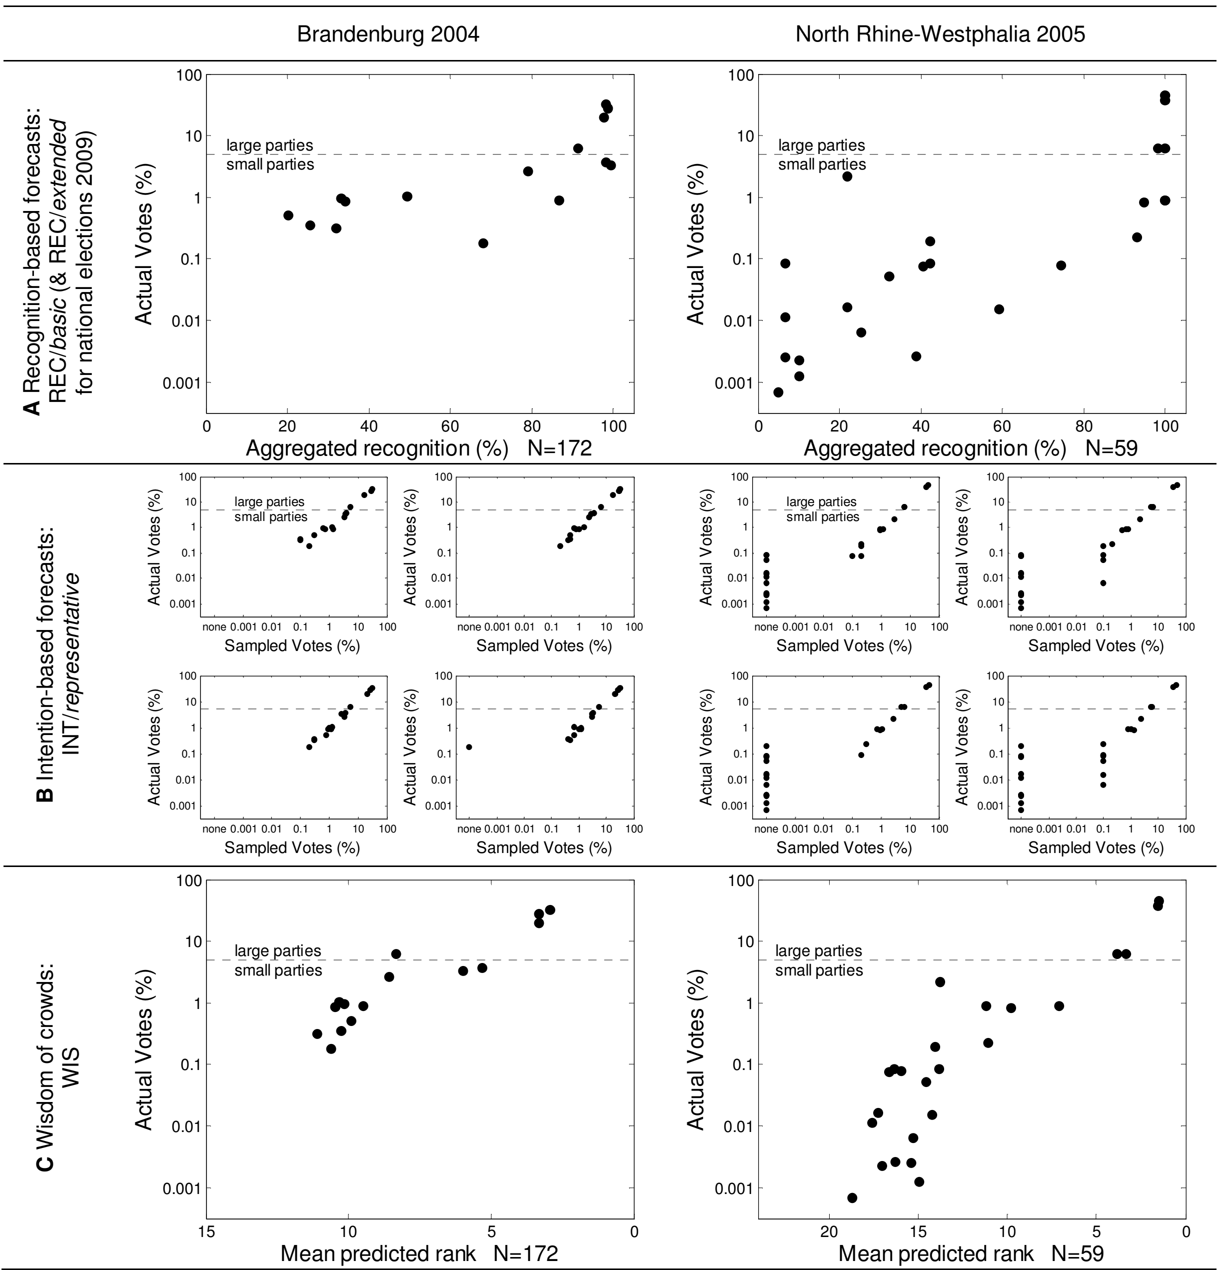

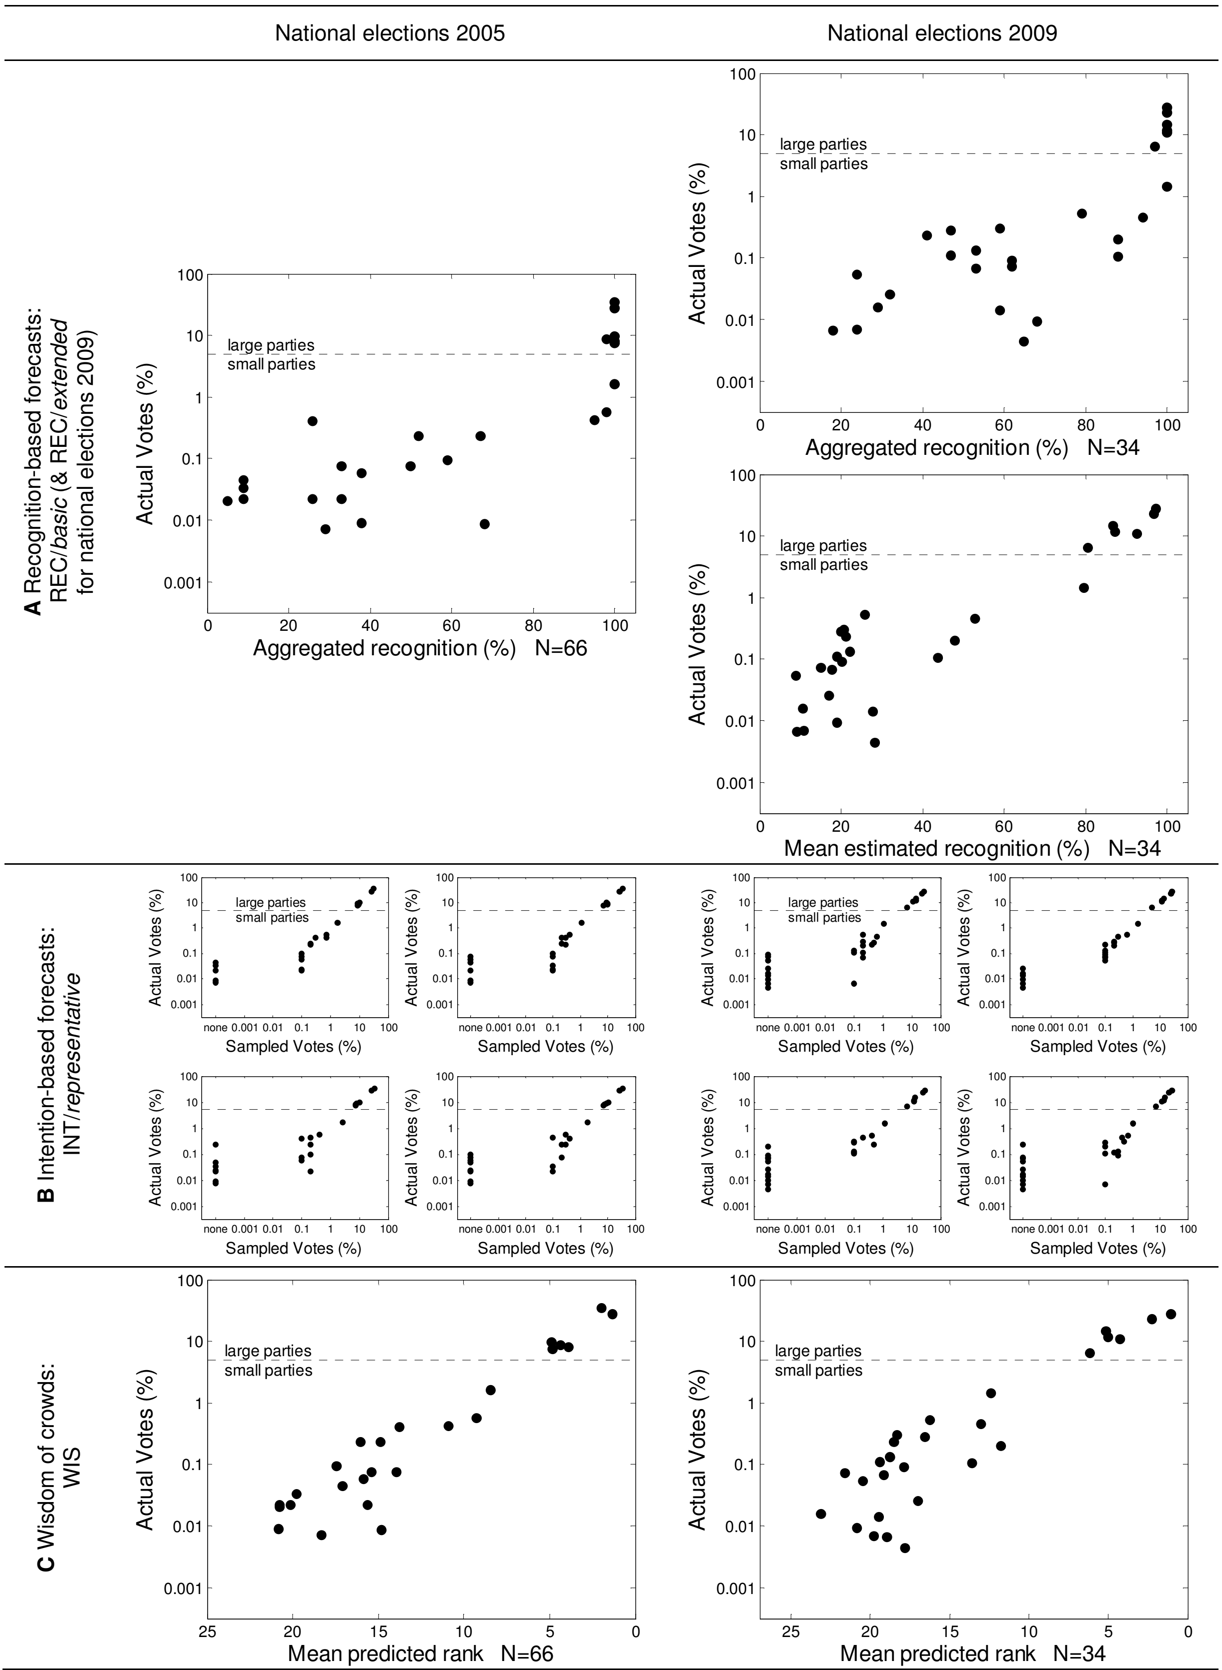

https://doi.org/10.1017/S1930297500002102 Published online by Cambridge University Press

that recognition is a robust predictor variable in election forecasts for smaller political parties.

It may seem somewhat counterintuitive that it is possible to forecast elections with such naïve, recognition-based methods, and in fact, we would like to point out that prior to conducting our first study in 2004, we did not expect recognition-based forecasts to perform as well as they did. As the first three studies represent reanalyses of already existing data, we retained our skepticism and thought it was particularly important to replicate these results in Study 4, in which we also added further competing models, such as REC/*extended*. Our results fit to a growing body of research showing that simple forecasting models perform often as good or even better as more complex ones (e.g., Brighton, 2006; Czerlinski, Gigerenzer, & Goldstein, 1999; Dawes, 1979; Einhorn & Hogarth, 1975; Gigerenzer & Brighton, 2009; Gigerenzer & Gaissmaier, 2011; Hogarth & Karelaia, 2007; Marewski, Gaissmaier, & Gigerenzer, 2010a, b). And indeed, recognition plays an important role in some of these simple models (e.g., Gigerenzer & Goldstein, 1996).

We hasten to add, however, that the usefulness of REC/*basic* for predicting elections is likely to be restricted to multi-party systems as they exist in many European countries. If only a few well-known parties compete (e.g., as Democrats and Republicans in the U.S.A), then the binary recognition judgments elicited in Studies 1–4 cannot discriminate between them and will not yield accurate predictions. At the same time, as we have pointed out above, even in multi-party systems the collective recognition used by REC/*basic* will not be a useful predictor variable for the larger political parties' electoral success, because these parties tend to be equally well recognized (see Figure 3). Furthermore, as suggested by the relative boost in performance of the intention-based forecasts in the Brandenburg election (Study 1), in which only 15 parties competed compared to 24 to 27 in the other elections, the relative usefulness of recognition-based forecasts in comparison to intention-based ones may be further limited to elections where many parties compete. Finally, recognition can be biased when parties are recognized for reasons unrelated to the parties' electoral success. This is likely the case for radical parties. To give just one example, consider Figure 3, panel A, for the national elections 2009: The party that actually received the lowest share of votes, 0.0044%, was the DKP ("German Communist Party"), yet this party was still recognized by about 65% of our participants.

Moreover, at the close of this article, we would like to stress that other simple forecasting methods may allow forecasting elections as accurately as or even more accurately than recognition. These methods include models that we did not test here, such as Lichtman's (2008) *keys model* or a version of the *take-the-best heuristic* (Graefe & Armstrong, in press), both of which were successful in forecasting presidential elections in the U.S.A.

In fact, also the other simple forecasting method that we actually did test—wisdom of crowds, WIS—was more successful than REC/*basic*'s forecasts, which echoes similar results in the literature demonstrating that wisdom-of-crowds-forecasts are quite accurate (e.g., Sjöberg, 2009).[9] In our studies, it is likely that WIS's success is fuelled by additional information the interviewed persons may have used to generate their individual predictions of the election outcomes, particularly to discriminate between two or more parties they recognized. This is most likely the case for the larger parties. These parties tend not only to be commonly recognized, but also people tend to know more about them than about the smaller parties; opinion polls and other information relevant for forecasting electoral success tend to be widely communicated by the media about these parties—not only prior to elections.

However, WIS also allowed better discriminating between the smaller parties than REC/*basic*. One explanation for this finding could be that some small parties are recognized by many people for reasons unrelated to electoral success, which holds true for extremely right-wing parties, for instance. If people are aware that they recognize a party name for reasons unrelated to electoral success, they may simply discount their recognition (Marewski, Gaissmaier, Schooler, et al., 2009; see also Oppenheimer, 2003, for similar findings in other domains). In principle, the party name could even allow people to discriminate between two unrecognized small parties, for instance when the party name is an absurd, satiric one (as in the eyes of many may be the case for the *Anarchistic Pogo Party*, although the authors do not take sides here). As a side note, Sjöberg (2009) actually speculated that knowledge of polls would be a major source for the success of wisdom of crowds, and in his case this may be true as he exclusively studied large parties. However, it is unlikely that polling results aided the per-

[9]To generate wisdom-of-crowds-forecasts, we asked people to predict the rank order of votes the parties would gain, and then averaged these rank orders, using the average ranks to forecast the election outcomes. As pointed out to us by Jon Baron and an anonymous reviewer, rather than averaging the rank predictions of election results across participants, it would have been interesting to ask participants for estimates of vote shares: "How many out of a 100 randomly drawn people do you think would vote for this party?" This would have allowed a more direct comparison with the extended recognition model REC/*extended* and would thus have helped telling whether averaged recognition estimates are largely fuelled by recognition per se, or by the fact that estimates (of some sort) are being aggregated. Unfortunately we did not collect corresponding data when we ran our studies. However, the WIS model that we tested is similar in principle, except that was based on predictions of ranks rather than shares.

formance of our wisdom-of-crowds-model WIS for the small parties we studied here, as such information is usually not available for these parties in Germany.

Finally, we wish to point out that even WIS did not outperform our second recognition-based forecasting model, REC/*extendend*, which bases forecasts on people's averaged estimates how many out of 100 randomly drawn people would recognize a party. These two models' performance was basically indistinguishable, suggesting that people seem to be able to successfully discriminate between highly recognizable parties (e.g., large parties, radical parties) when estimating population recognition rates. In fact, as much as it is possible that people base the election forecasts used in WIS on recognition (see above), it is *also* possible that people's estimates of other people's recognition are at least partially based on the same information that may come to bear in WIS: For instance, if a person knows she recognizes the party "Grey Panthers"—a small party for the elderly—exclusively because her grandmother happens to be a member of this party, then the person may discount her recognition of this party name and adjust her estimate of the population recognition rate accordingly.

Let us conclude by returning to the dilemma faced by Leonid Brezhnev, who, as pointed out in the beginning, once remarked that "The trouble with free elections is, you never know who is going to win" (Rees, 2006). Brezhnev's dilemma can be solved in various ways: abolishing free elections, manipulating who will win, or relying on surveys of voting intentions to find out who will win in advance. We have contributed to develop yet another solution. As we have shown, simple forecasting models based on collective recognition, people's estimates of other people's recognition, or the aggregated wisdom of many may help forecasting who will win. Admittedly, this may not be the solution that Brezhnev had in mind.

References

Abramowitz, A. I. (1975). Name familiarity, reputation, and the incumbency effect in a congressional election. *Western Political Quarterly, 27*, 668–684.

Armstrong, J. S. (2001). Combining forecasts. In J. S. Armstrong (Ed.), *Principles of forecasting: A handbook for researchers and practitioners* (pp. 417–439). Norwell, MA: Kluwer Academic.

Borges, B., Goldstein, D. G., Ortmann, A., & Gigerenzer, G. (1999). Can ignorance beat the stock market? In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds), *Simple heuristics that make us smart* (pp. 59–72). New York: Oxford University Press.

Brighton, H. (2006). Robust inference with simple cognitive models. In C. Lebiere & B. Wray (Eds.), *Between a rock and a hard place: Cognitive science principles meet AI-hard problems. Papers from the AAAI Spring Symposium* (AAAI Tech. Rep. No. SS-06–03, pp. 17–22). Menlo Park, CA: AAAI Press.

Campbell, J. E., Alford, J. R., & Henry, K. (1984). Television Markets and Congressional Elections. *Legislative Studies Quarterly, 9*, 665–678.

Campbell, J. E., & Lewis-Beck, M. S. (2008). US presidential election forecasting: An introduction. *International Journal of Forecasting, 24*, 189–192.

Clemen, R. T. (1989). Combining forecasts: A review and annotated bibliography. *International Journal of Forecasting, 5*, 559–583.

Coates, S. L., Butler, L. T., & Berry, D. C. (2004). Implicit memory: A prime example for brand consideration and choice. *Applied Cognitive Psychology, 18*, 1195–1211.

Coates, S. L., Butler, L. T., & Berry, D. C. (2006). Implicit memory and consumer choice: The mediating role of brand familiarity. *Applied Cognitive Psychology, 20*, 1101–1116.

Craik, F. I. M. & McDowd, M. (1987). Age differences in recall and recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 474–479.

Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer, P. M. Todd, & the ABC Research Group, *Simple heuristics that make us smart* (pp. 97–118). New York: Oxford University Press.

Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist, 34*, 571–582.

Einhorn, H. J., & Hogarth, R. M. (1975). Unit weighting schemes for decision making. *Organizational Behavior and Human Performance, 13*, 171–192.

Galton, F. (1907). Vox populi. *Nature. 75*, 7.

Gigerenzer, G. (1982). Der eindiminsionale Wähler. *Zeitschrift für Sozialpsychologie*, *13*, 217–236.

Gigerenzer, G. (2007). *Gut feelings. The intelligence of the unconscious*. New York: Viking.

Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science, 1*, 107–143.

Gigerenzer, G. & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology, 62*, 451–482.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *104*, 650–669.

Goldenberg, E. N. & Traugott, M. W. (1980). Congressional campaign effects on candidate recognition and evaluation. *Political Behavior, 2*, 61–90.

Goldstein, D. G. (2007). Getting attention for unrecognized brands. *Harvard Business Review*, *85*, 24–28.

Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, *109*, 75–90.

Goodman, L. A., & Kruskal, W. H. (1954). Measures of association for cross classifications. *Journal of the American Statistical Association*, *49*, 732–764.

Graefe, A., & Armstrong, J. S. (in press). Predicting elections from the most important issue: A test of the take-the-best heuristic. *Journal of Behavioral Decision Making.*

Hertwig, R., & Todd, P. M. (2003). More is not always better: The benefits of cognitive limits. In D. Hardman & L. Macchi (Eds.), *Thinking: Psychological perspectives on reasoning, judgment and decision making* (pp. 213–231). Chichester, England: Wiley.

Hertwig, R., Herzog, S., Schooler, L. J., & Reimer, T. (2008). Fluency heuristic: A model of how the mind exploits a by-product of information retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*, 1191–1206.

Herzog, S. M., & Hertwig, R. (2011). The wisdom of ignorant crowds: Predicting sport outcomes by mere recognition. *Judgment and Decision Making*.

Hilbig, B. E. (2008). Individual differences in fast-and-frugal decision making: neuroticism and the recognition heuristic. *Journal of Research in Personality, 42*, 1641–1645.

Hilbig, B. E. & Pohl, R. F. (2008). Recognizing users of the recognition heuristic. *Experimental Psychology, 55*, 394–401.

Hogarth, R. M. (1978). A note on aggregating opinions. *Organizational Behavior and Human Performance*, *21*, 40–46.

Hogarth, R. M., & Karelaia, N. (2007). Heuristic and linear models of judgment: Matching rules and environments. *Psychological Review*, *114*, 733–758.

Jackman, S., & Sniderman, P. M. (2002). The institutional organization of choice spaces: A political conception of political psychology. In K. Monroe (Ed.), *Political psychology* (pp. 209–224). Mahway, NJ: Erlbaum.

Jacobson, G. C. (1987). *The Political of Congressional Elections*. Boston: Little, Brown.

Johnson, T. R., Budescu, D. V., & Wallsten, T. S. (2001). Averaging probability judgments: Monte Carlo analyses of asymptotic diagnostic value. *Journal of Behavioral Decision Making*, *14*, 123–140.

Kelley, S., & Mirer, T. W. (1974). The simple act of voting. *American Political Science Review*, *68*, 572–591.

Lewis-Beck, M. S., & Rice, T. W. (1992). *Forecasting Elections.* Washington, D. C.: Congressional Quarterly Press

Lichtman, Allan J. (2008), The keys to the white house: An index forecast for 2008. *International Journal of Forecasting*, *24*, 301–309.

Mann, T. E., & Wolfinger, R. E. (1980). Candidates and Parties in Congressional Elections. *American Political Science Review*, *74*, 617–632.

Marewski, J. N., Gaissmaier, W., & Gigerenzer, G. (2010a). Good judgments do not require complex cognition. *Cognitive Processing*, *11*, 103–121.

Marewski, J. N., Gaissmaier, W., & Gigerenzer, G. (2010b). We favor formal models of heuristics rather than lists of loose dichotomies: a reply to Evans and Over. *Cognitive Processing, 11*, 177–179.

Marewski, J. N., Gaissmaier, W., Schooler, L. J., Goldstein, D. G., & Gigerenzer, G. (2009). Do voters use episodic knowledge to rely on recognition? In N. A. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 2232–2237). Austin, TX: Cognitive Science Society.

Marewski, J. N., Gaissmaier, W., Schooler, L. J., Goldstein, D. G., & Gigerenzer, G. (2010). From Recognition to Decisions: Extending and Testing Recognition-Based Models for Multi-Alternative Inference. *Psychonomic Bulletin and Review*, *17*, 287- 309.

Miller, J. M., & Krosnick, J. A. (1998). The impact of candidate name order on election outcomes. *Public Opinion Quarterly*, *62*, 291–330.

Newell, B. R., & Fernandez, D. (2006). On the binary quality of recognition and the inconsequentiality of further knowledge: Two critical tests of the recognition heuristic. *Journal of Behavioral Decision Making*, *19*, 333–346.

Oppenheimer, D. M. (2003). Not so fast! (and not so frugal!): Rethinking the recognition heuristic. *Cognition*, *90*, B1–B9.

Pachur, T., & Biele, G. (2007). Forecasting from ignorance: The use and usefulness of recognition in lay predictions of sports events. *Acta Psychologica*, *125*, 99–116.

Pachur, T., Bröder, A., & Marewski, J. N. (2008). The recognition heuristic in memory-based inference: Is recognition a non-compensatory cue? *Journal of Behavioral Decision Making*, *21*, 183–210.

Pachur, T., & Hertwig, R. (2006). On the psychology of the recognition heuristic: Retrieval primacy as a key determinant of its use. *Journal of Experimental Psy-*

chology: Learning, Memory, and Cognition, *32*, 983–1002.

Pachur, T., Todd, P. M., Gigerenzer, G., Schooler, L. J., & Goldstein, D. G. (in press). Is ignorance an adaptive tool? A review of recognition heuristic research. In P. M. Todd, G. Gigerenzer, & the ABC Research Group, *Ecological rationality: Intelligence in the world*. New York: Oxford University Press.

Pohl, R. (2006). Empirical tests of the recognition heuristic. *Journal of Behavioral Decision Making*, *19*, 251–271.

Popkin, S. L. (1994). *The reasoning voter: Communication and persuasion in presidential campaigns*, 2d edition. Chicago: University Of Chicago Press.

Rees, N. (2006). *Brewer's famous quotations: 5000 quotations and the stories behind them.* London: Weidenfeld & Nicolson.

Reimer, T., & Katsikopoulos, K. V. (2004). The use of recognition in group decision-making. *Cognitive Science*, *28*, 1009–1029.

Scheibehenne, B., & Bröder, A. (2007). Can lay people be as accurate as experts in predicting the results of Wimbledon 2005? *International Journal of Forecasting*, *23*, 415–426.

Schooler, L. J., & Hertwig, R. (2005). How forgetting aids heuristic inference. *Psychological Review, 112,* 610–628.

Serwe, S., & Frings, C. (2006). Who will win Wimbledon? The recognition heuristic in predicting sports events. *Journal of Behavioral Decision Making, 19,* 321–332.

Shepard, R. N. (1967). Recognition memory for words, sentences, and pictures. *Journal of Verbal Learning and Verbal Behavior, 6,* 156 –163.

Sigelman, L., Batchelor, R., & Stekler, H. (1999). Political Forecasting. *International Journal of Forecasting*, *15*, 125−126.

Sjöberg, L. (2009). Are all crowds equally wise? A comparison of political election forecasts by experts and the public. *Journal of Forecasting, 28,* 1–18.

Sniderman, P. M. (2000). Taking sides: A fixed choice theory of political reasoning. In A. Lupia, M. D. McCubbins, & S. L. Popkin (Eds.), *Elements of reason: Cognition, choice, and the bounds of rationality* (pp. 74–84). New York: Cambridge University Press.

Standing, L. (1973). Learning 10,000 pictures. *Quarterly Journal of Experimental Psychology*, *25*, 207–222.

Stone, W. J., & Abramowitz, A. I. (1983). Winning may not be everything, but it's more than we thought: Presidential party activists in 1980. *American Political Science Review*, *77*, 945–956.

Surowiecki, J. (2004). *The wisdom of crowds.* New York: Doubleday.

Timmermann, A. (2006). Forecast combinations. In G. Elliott, C. Granger, & A. Timmermann (Eds.), *Handbook of economic forecasting* (pp. 135–196). Amsterdam: North Holland.

Todorov, A., Mandisodza, A. N., Goren, A., & Hall, C. (2005). Inference of competence from faces predict election outcomes. *Science*, *308*, 1623–1626.

Volz, K. G., Schooler, L. J., Schubotz, R. I., Raab, M., Gigerenzer, G., & Cramon, D. Y. von. (2006). Why you think Milan is larger than Modena: Neural correlates of the recognition heuristic. *Journal of Cognitive Neuroscience*, *18*, 1924–1936.

Wang, X. T. (2008). Decision heuristics as predictors of public choice. *Journal of Behavioral Decision Making*, *21*, 77–89.

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

Zajonc, R. B. (1968). Attitudinal effects of mere exposures. *Journal of Personality and Social Psychology*, *9*, 1–27.