**High Chances and Close Margins:
How Equivalent Forecasts Yield Different Beliefs**

Oleg Urminsky (oleg.urminsky@chicagobooth.edu)

University of Chicago

5807 S. Woodlawn Ave, Chicago, IL 60637 USA

Luxi Shen (luxi.shen@cuhk.edu.hk)

Chinese University of Hong Kong

12 Chak Cheung Street, Hong Kong

First Draft: 2/8/2018

Current Draft: 1/2/2019

Please contact authors for current version before citing.

Abstract

Statistical forecasts are increasingly prevalent. How do forecasts affect people's beliefs about corresponding future events? This research proposes that the format in which the forecast is communicated biases its interpretation. We contrast two common forecast formats: *chance* (e.g., the forecasted probability that a political candidate or a sports team will win) versus *margin* (e.g., by how many points a political candidate or a sports team is forecasted to win). Across five studies (total N=2,486; plus nine replications with an additional total N=2,420), we find a robust chance-margin discrepancy: chance forecasts lead to more extreme beliefs about outcome occurrence than margin forecasts. This discrepancy persists over time in interpreting real-world forecasts (e.g., the 2016 U.S. presidential election), replicates even when the forecasts are strictly statistically equivalent, and has downstream consequences for attitudes toward election candidates and sports betting decisions. Lastly, we discuss the implications for mass communications and gerrymandered redistricting.

Keywords: Judgment; Decision Making; Numerical Cognition; Forecasting; Intuitive Statistics; Political Beliefs

Forecasts are pervasive. People frequently encounter forecasts about weather, stock market performance, sports outcomes, and election results. Election forecasts reported in the mass media, for example, contribute to voters' beliefs about important events such as the United States presidential election and the Brexit referendum in the United Kingdom. In the summer of 2016, the media described the same status of pre-election opinion with headlines such as "Hillary Clinton Maintains 5-Point Lead Over Donald Trump" in the Wall Street Journal (Hook, 2016) and "The Upshot's Election Forecast: 76% Chance of a Clinton Victory" in the New York Times (Katz, 2016). These headlines framed forecasts in different ways for news readers, directing the readers' attention to either the predicted difference in the votes received by the candidates or the probability of a candidate winning. Thus, it is critical for us to understand how lay people comprehend forecasts and, what may bias their comprehension in which direction.

In this research, we test whether people's beliefs about a future outcome are shaped by the format of a forecast, beyond the content of the forecast, and find a chance-margin discrepancy: chance forecasts lead to more extreme beliefs about outcome occurrence than margin forecasts. For example, in an election, people see a candidate with a 64% predicted chance of winning as having a greater lead over the opponent than if the candidate is forecasted to win 51.5% of the votes, even though both forecasts are based on the same polling data and supposed to be equivalent.

In this article, we first discuss how forecasts of future events are generated and communicated, and then review literature on how lay people may interpret forecasts. Across five studies (total $N=2,486$; plus nine replications with an additional total $N=2,420$), we investigate a robust chance-margin discrepancy, such that chance forecasts lead to more extreme beliefs about outcome occurrence than margin forecasts. Our empirical evidence reveals that this discrepancy occurs due to failed statistical reasoning, persists over time in interpreting real-world forecasts (e.g.,

the 2016 U.S. presidential election), replicates even when the forecasts are strictly statistically equivalent, changes consequential decisions (e.g., sports betting) and is not explained by simple heuristics, such as anchoring or number matching. Our research makes unique contributions to the literature on judgement and decision making in general and specifically to research on statistical reasoning, and has implications for media and policy communication of forecasts.

Since the early days of statistics, experts have used different approaches to generate and communicate forecasts about future events. In the past, pundits would use their personal judgment or exemplars (e.g., “bellwether” locations) to make predictions about election outcomes. Nowadays, election forecasters have increasingly been taking a more scientific approach, relying on political polling results and model-based polling aggregation (Hilygus, 2011). These modern approaches first collect data (e.g., from political polling) and then use the data to generate a full, estimated probability distribution of the potential outcomes. However, these distribution-level predictive statistics are rarely conveyed to the mass audience. Instead, lay people typically encounter a single statistic, computed from the full probability distribution and used to summarize that distribution, which is presented as the forecast.

Figure 1: A chance forecast and a margin forecast that are based on the same prediction distribution

These forecasts are commonly communicated in one of two formats: (a) the *chances* of a focal outcome (e.g., the predicted probability that a political candidate will win the election, or the probability that a sports team will win the game) or (b) the expected *margin* (e.g., the share of votes a candidate is expected to receive over their opponent in an election, or the expected point-spread between two teams). Since both forecasts are based on the same underlying probability distribution (Figure 1), they summarize the same outcome distribution, and for a given distribution (e.g., a normal distribution with known variance), the forecasts are statistically equivalent. However, are they also equivalent psychologically? The extant literature yields diverging evidence about the robustness and accuracy of how people process statistical information, suggesting conflicting answers to our research question.

One possibility is that people do see through the format difference and process chance and margin forecasts as if they were statistically equivalent. This possibility is theoretically grounded in

the view, prevalent in cognitive and mathematical psychology, that people are skilled intuitive statisticians, particularly in domains in which they have repeated exposure and ample experience, such as politics and sports. Recent influential research in cognitive psychology has argued that many decision-making phenomena formerly regarded as errors may actually be explained with Bayesian models that account for uncertainty or the computational cost of complex inferences (Lieder, Griffiths, & Goodman, 2012; Pouget, Beck, Ma, & Latham, 2013).

According to this view, people have a nearly-optimal ability to represent and update probability distribution information (Griffiths & Tenenbaum, 2006) that perhaps is even spontaneous at the neural level (Knill & Pouget, 2004; Ma, Beck, Latham, & Pouget, 2006). If this is indeed the case, people should be able to correctly interpret corresponding chance and margin forecasts from the same underlying outcome distribution, thereby extracting equivalent information from the different forecast formats. However, a recent debate (Bowers & Davis, 2012; Griffiths, Chater, Norris, & Pouget, 2012; Jones & Love, 2011) has questioned the degree to which normative Bayesian processes can fully explain performance in complex reasoning and decision tasks.

Another possibility is that numerical information, including forecasts, may be processed differently when communicated in a different format. This prediction stems from the contrary view, that people are poor intuitive statisticians and are often systematically biased in their statistical reasoning. A classic literature on intuitive statistics suggests that this could be the case, particularly when people make inferences from samples (Kahneman & Tversky, 1972; Peterson, DuCharme, & Edwards, 1968; Wheeler & Beach, 1968). In particular, people have been shown to form different interpretations when equivalent information is framed differently (e.g., Levin & Gaeth, 1988; Kuhberger, 1998; Tversky & Kahneman, 1981). Recent work on numerical cognition has further demonstrated important consequences of information format on beliefs and decision-making.

Larrick and Soll (2008) found that although automobile buyers care about gas consumption and carbon emissions, they nevertheless misunderstand “mile per gas (MPG)” information, incorrectly assuming that the amount of gas consumed by a car decreases as a linear function of the car’s MPG. Thus, people may fall short of the Bayesian ideal, relying on an incomplete understanding to process forecasts, resulting in miscalibrated interpretations.

Next, we present five studies that test whether people interpret chance and margin forecasts equivalently in real-world forecasting contexts (elections, Study 1), and accurately use equivalent chance forecasts or margin forecasts they have been presented with to estimate the unknown forecast (elections, Study 2; sports outcomes, Study 3). We then investigate how forecast format affects incentive-compatible choices (sports betting, Study 4). Lastly, as a strong test of internal validity, we elicited incentive-compatible estimation responses to forecasts in a statistical scenario that ensures the strict mathematical equivalence of the chance and margin forecasts (Study 5). We provide additional details and replication studies in Appendix, and all experimental details and data are available at OSF.

Study 1: Forecast Format Impacts Subjective Assessments

During the 2016 U.S. presidential election, we tested whether people interpreted equivalent chance and margin forecasts equivalently. We presented participants with either a chance forecast (probability of election outcome) or an ostensibly equivalent margin forecast (predicted vote share) and asked participants to rate their attitude towards the news conveyed by the forecast. If people’s beliefs are only shaped by the content of the forecast regardless of the format, then they should have the same reaction to both forecasts. However, if people interpret chance forecasts as conveying a

stronger lead by one candidate than what is conveyed by margin forecasts, then we would expect to see more extreme attitudinal responses to the chance forecast.

Method

In this study ($N = 225$ valid AMT participants), participants saw an election forecast, either displaying the chances of Clinton winning or displaying Clinton's predicted margin of victory. This study was conducted four days before the election, when Clinton was leading Trump 51.5% vs. 48.5%, and Clinton was projected to have a 64% chance of winning, according to *fivethirtyeight.com*.

Participants were asked to assess their opinions of the current predictions on a scale from 1 (“*very good news*”) to 7 (“*very bad news*”), along with several other questions (see Appendix). Since a wider perceived lead for Clinton would be interpreted as better news by her supporters but worse news by Trump supporters, we coded the extremity of rated attitudes as the absolute difference from the midpoint of the scale ($\text{extremity} = |\text{rating} - 4|$).

Results

As predicted, participants in the chance-forecast-display condition had more extreme reactions to the state of the election conveyed by the forecast, than those in the margin-forecast-display condition ($M_{\text{chance-forecast-display}} = 1.76$ vs. $M_{\text{margin-forecast-display}} = 1.38$, $d = 0.32$, $t(223) = 2.41$, $p = .017$). This demonstrates an effect of forecast format on attitudes, such that equivalent forecasts nevertheless result in diverging attitudes. Notably, this effect on attitudes cannot be explained by anchoring-and-adjustment, as the 7-point scale used by participants was different than the forecast scale (0-100%) and did not include the forecast values. We replicated this finding in an additional study (see Appendix) where people were presented with information about changes in forecasts.

Study 2: Presidential Election Forecasts of Winning Chances vs. Vote Margin

Study 1 demonstrates that ostensible equivalent forecasts communicated in different formats yield different attitudes about the election. We believe that this occurs because lay people do not interpret the forecasts correctly. We propose that people's intuitions about chance and margin forecasts for election outcomes are inaccurate, diverging from the real information conveyed by forecasts in a systematic way. We suspect that a systematic bias in estimation occurs because of a lack of understanding about the non-linear relationship between the forecast formats (Figure 1). As a result, lay people would perceive chance forecasts as more extreme than ostensibly equivalent margin forecasts from the same predicted distribution. If this is the case, people would overestimate margins when shown chance forecasts and underestimate chance forecasts when shown margin forecasts, which we test in the next study.

Method

We tested whether people interpreted equivalent chance and margin forecasts accurately over the course of the 2016 U.S. presidential election. We presented participants with either a chance forecast (probability of election outcome) or an ostensibly equivalent margin forecast (predicted vote share) and asked participants to estimate the forecast in the other format. If people's beliefs are only shaped by the content of the forecast, regardless of the format, then they should be able to use one forecast to estimate the equivalent forecast in the alternative format, without systematic errors.

We collected data in six waves, between August and November 2016, prior to the 2016 election. A combined 1,163 Amazon Mechanical Turk (AMT) workers completed surveys, after

excluding those with duplicate IP addresses and failed attention checks (which was done in all the studies). We used actual election forecasts from the website *fivethirtyeight.com*, based on aggregated political opinion polls, of (1) the chance (probability) of each candidate winning and (2) the predicted vote-margin for each candidate. Both forecasts were generated from the same statistical model, based on the same polling data, and therefore summarized the same predicted distribution of outcomes. As the results of the election polls used to generate the predicted distribution varied over time, the forecasts varied as well.

Participants were randomly assigned to one of two conditions. In the margin-forecast-display condition, they saw the predicted vote-margin for both candidates and were asked to estimate the corresponding chance forecast for each candidate. In the chance-forecast-display condition, they saw the probability of each candidate winning and were asked to estimate the corresponding margin forecast for each candidate. Importantly, they were not asked about their own personal predictions, but rather were asked to estimate the predictions made by the website. We counterbalanced the order in which the candidates were presented in all waves except the first, and the order of presentation did not affect the results.

Results

Participants' estimates were significantly biased (see violin plots in Figure 2), suggesting that they did not interpret the two forecasts formats as representing the same underlying political situation. Overall, participants in the chance-forecast-display condition, who saw a chance forecast (an average 74.2% chance of Clinton winning), overestimated the margin forecast (60.5% estimated share for Clinton vs. 52.6% actual share, $d = 0.58$, $t(575) = 13.91$, $p < .001$). Participants in the margin-forecast-display condition, who were shown a margin forecast (an average share for Clinton

of 52.6%), underestimated the chance forecast (59.5% estimated chance of Clinton winning vs. 74.2% actual chance, $d = 0.87$, $t(586) = 21.16$, $p < .001$). Both forecast formats yielded significantly biased estimates, but the estimates were biased in opposite directions—while participants underestimated the chance forecasts, they overestimated the margin forecasts. These findings replicated in each of the six waves, as the actual forecasts varied (all $ps < .001$).

Figure 2: Estimation results based on different forecast formats in Study 1

Black diamonds represent the actual election forecast values at the time of the survey. Colored dots indicate the mean estimates of one format of forecast, made by participants who had seen the other forecast format, with an error bar representing the confidence interval. The colored plot shapes display the distribution of estimates. Estimated chance forecasts are shown in orange and estimated margin forecasts are shown in blue.

A regression predicting participants' margin forecast estimates in the chance-forecast-display condition found significant sensitivity to the true margin forecast value ($\beta_{\text{TRUE}} = 2.18$, $\eta^2 = 0.032$, $t(574) = 4.35$; $p < .001$). Likewise, a regression predicting participants' chance forecast estimates in the margin-forecast-display condition found significant sensitivity to the true chance forecast value ($\beta_{\text{TRUE}} = 0.26$, $\eta^2 = 0.034$, $t(585) = 4.54$; $p < .001$). We also found a significant interaction between estimation condition and true value ($\beta_{\text{FORMAT} \times \text{TRUE}} = -0.88$, partial $\eta^2 = 0.099$, $t(1159) = 11.34$; $p < .001$) in a follow-up regression predicting estimation error. These results suggest that while participant estimates were sensitive to the changes in the state of the election over time, there was a robust forecast format bias that was larger (more underestimation in the margin-forecast display condition and more overestimation in the chance-forecast display condition) the more extreme the actual forecasts. The vast majority of participants displayed the chance-margin discrepancy—80% of participants in the chance-forecast-display condition overestimated the margin while 82% of participants in the margin-forecast-display condition underestimated the chance, and these results were not explained by participants reporting back the number they had been shown.

Study 3: Sports Forecasts of Winning Chances vs. Point-Spreads

Our findings about how forecast format impacts people's interpretation of election forecasts rely on the assumption that the two summary statistics (chance and margin) are in fact equivalent and represent the same accurate prediction distribution. During the election, however, different forecasters disagreed about how to translate the polling-based forecasted vote share into the candidates' chances of winning the election. In particular, the distribution of votes across states and

the role of the Electoral College complicated chance predictions. We selected *fivethirtyeight.com* because their chance predictions were the most conservative (e.g., compared to the Princeton Election Consortium). However, the 2016 U.S. presidential election outcome did not turn out as predicted and evidence suggests that polling aggregation was non-optimal (Wright & Wright, 2018). Thus, we cannot exclude the possibility—however unlikely—that lay people’s estimates in our studies, averaged together, were in fact more accurate than the professional forecasts. To assess generality, we therefore test the effect of forecast format in sports forecasts, a domain different from political elections.

Method

In a pre-registered Study 3 (AsPredicted #6391, N = 258 valid AMT participants), we compared inferences about NBA basketball games when people were shown either the predicted chances of winning or the predicted margin of winning (point-spread) and then estimated the other forecast format, as in Study 1. We used chance and point-spread forecasts from *fivethirtyeight.com* for four basketball games that took place in October 2017, shortly before the games were played. These kinds of sports forecasts have been demonstrated to be highly accurate (Spann & Skiera, 2009).

The matches ranged from almost even, with neither team strongly favored (Rockets vs. Hornets, 54% vs. 46%, 1 point), to very uneven, with one team highly favored (Warriors vs. Wizards, 84% vs. 16%, 10 points spread). The two other games lay in between (Nuggets vs. Hawks, 70% vs. 30%, 5.5 points; Knicks vs. Nets, 59% vs. 41%, 2.5 points). Unlike the election results, each of the four predictions turned out to be correct.¹ Each participant saw either the forecasted chances of

¹ Retrieved from https://www.basketball-reference.com/leagues/NBA_2018_games.html

winning for both teams in each of the four games and estimated the corresponding point-spreads (in the chance-forecast-display condition) or the point-spread for each game and estimated the corresponding forecasted chances of winning (in the margin-forecast-display condition). An additional advantage of this context is that point-spread is measured on a completely different scale (number of points) from chances (percentage). We counterbalanced the order in which the teams were presented and also varied the framing (e.g., Warriors winning vs. Wizards losing), and neither affected the results. As in Study 1, participants were instructed to estimate what the website predicted, rather than state their own beliefs.

Results

Participants in the chance-forecast-display condition overestimated the point-spread ($p < .001$ for all four games, left panel of Figure 3), with an average overestimate of +8.73 (13.48 estimated vs 4.75 actual, $d = 0.94$, $t(102) = 9.49$, $p < .001$). By contrast, participants in the margin-forecast-display condition significantly underestimated the forecasted chances of winning ($ps < .001$) except for the Rockets-Hornets game ($p = .517$), with an average underestimate of -8.35 percentage points (58.40 estimated vs 66.75 actual, $d = 1.11$, $t(117) = 12.11$, $p < .001$).

A regression with clustered standard errors predicting participants' margin forecast estimates in the chance-forecast-display condition found significant sensitivity to the true margin forecast value ($\beta_{\text{TRUE}} = 1.66$, $\eta^2 = 0.21$, $t(410) = 11.51$, $p < .001$). Likewise, a regression predicting participants' chance forecast estimates in the margin-forecast-display condition found significant sensitivity to the true chance forecast value ($\beta_{\text{TRUE}} = 0.44$, $\eta^2 = 0.10$, $t(470) = 7.15$, $p < .001$). We again found a significant interaction between estimation condition and true value ($\beta_{\text{FORMAT} \times \text{TRUE}} = -0.76$, partial $\eta^2 = 0.097$, $t(880) = 10.12$, $p < .001$) in a follow-up regression predicting estimation

error. These results suggest that while participant estimates were sensitive to the differences in odds of winning over the games, there was a robust forecast format bias that was larger the more extreme the actual forecasts.

Figure 3: Estimation results based on different forecast formats in Study 3

Black diamonds represent the actual forecasts for each NBA game. Colored dots indicate the mean estimates of one format of forecast, made by participants who had seen the other forecast format, with an error bar representing the confidence interval. The colored plot shapes display the distribution of estimates. Estimated chance forecasts are shown in orange and estimated margin forecasts are shown in blue.

These results extend our findings of the chance-margin discrepancy from political elections to NBA games. Participants in the chance-forecast-display condition overestimated point-spreads, while participants in the margin-forecast-display condition (i.e. shown points-spread forecasts) underestimated chances of winning. This bias was pervasive across participants—93% of participants in the chance-forecast-display condition overestimated the point-spread, while 88% of participants in the margin-forecast-display condition underestimated the chance forecast, and this was not explained by participants simply repeating the forecast shown. The chance-margin discrepancy also was not explained by lack of familiarity with the context, as it was not moderated by level of involvement with NBA games.

Study 4: Impact of Forecast Format on Betting

Study 3 demonstrated the effect of forecast format on beliefs about NBA games. Could such forecast-induced beliefs change people's behavior? Study 4 tests whether forecast format impacts people's willingness to bet on the team predicted to win an upcoming game.

Method

In the pre-registered Study 4 (AsPredicted #8906, N = 387 valid AMT participants), participants read about an upcoming NBA game in which one team was strongly favored to win (Toronto Raptors with 93% chance vs. Atlanta Hawks with 7% chance; Toronto Raptors to win with a point-spread of 16). Participants saw the team names and either the point-spread (in the margin-forecast-display condition) or the forecasted chances (in the chance-forecast-display condition) for both teams. Participants were told that five lottery winners would receive a \$5 bonus each.

Before finding out if they had won, they could choose to bet as much of the potential \$5 as they wished on the team more likely to win (Raptors) and keep the remainder. Their bets would pay off \$3 for each \$1 bet on the winning team, but they would lose the money if they bet on the losing team. Participants then allocated the \$5 between betting on the team predicted to win and not betting. After the study had completed, the lottery winners were notified and paid based on their bets.

Results

Participants bet more in the chance-forecast-display condition than in the margin-forecast-display condition, when shown the point-spread, ($M_{\text{chance-forecast-display}} = \3.03 vs. $M_{\text{margin-forecast-display}} = \2.62 , $d = .22$, $t(369) = 2.15$, $p = .033$). While the difference in the likelihood of betting any money was not significant (80% vs. 84%, $\chi^2 = .78$, Fisher $p = .35$), those participants who did bet on the Raptors (rather than keep all their money) chose to wager more money after seeing the chance forecast ($M_{\text{chance-forecast-display}} = \3.61 vs. $M_{\text{margin-forecast-display}} = \3.28 , $d = 0.23$, $t(302) = 2.01$, $p = .045$). The effect of forecast format on betting did not differ by level of basketball experience (knowledge about and interest in NBA basketball or number of games watched). This suggests that the impact of format on betting was not due to a lack of relevant knowledge. Rather, a chance (vs. margin) forecast led to higher confidence about the outcome, and this higher confidence was expressed by the willingness to bet more on the game.

Study 5: Accuracy of Estimates Based on Forecast Format in a Statistical Scenario

The findings of the studies thus far suggest that people treat ostensibly equivalent forecasts differently as if chance forecasts (in both elections and sports) convey a more polarized reality than

margin forecasts. However, this assumes that the chance and margin forecasts are generated from a full outcome distribution that is both correct and available to the participants. In both the election and sports contexts, this assumption may fail in either regard: the actual shape of the full distribution may differ from the forecasters' assumptions, or the distribution might not be knowable to participants.

In this study, we provide people with sufficient objective information to enable a Bayesian presented with the chance forecast to generate the exact margin forecast, and vice versa. We then test the effects of forecast format on people's judgments in an incentive-compatible statistical scenario.

Method

In the pre-registered Study 5 (AsPredicted #6686, $N = 245$ valid AMT participants), participants read a statistical scenario involving a jar containing 99 mixed red and green marbles with the number of red marbles chosen at random from a uniform distribution. Four independent samples were generated by drawing 20 marbles from the jar with replacement and recording the number of red marbles for each sample. The samples were not observed by the participants. However, a statistician accurately determined, based on each sample, both (1) the chance that there were more red than green marbles in the jar and (2) the expected margin (i.e., how many more red marbles than green marbles were expected).

To calculate the forecasts, we applied Bayes Theorem (see Appendix for details). We can think of the margin forecast (i.e., expected number of red vs. green marbles), based on a given sample drawn with replacement, as the expected value of the posterior distribution. We can think of the

probability of a majority of marbles being red, conditional on the sample outcome, as the mass of the posterior distribution above 50% red.

Participants were randomly assigned to one of two conditions. In the chance-forecast-display condition, they were shown the chance forecast (the probability of more red than green marbles) for each of the four samples and estimated the corresponding margin forecasts. In the margin-forecast-display condition, participants were shown the margin forecasts (how many more red than green marbles were expected) for each of the four samples and then estimated the corresponding chance forecasts. We counterbalanced the order in which we presented the forecast information for red and green marbles. Each participant was paid a linear incentive for the accuracy of their estimates, with a maximum accuracy bonus of \$1 per person.

Results

As shown in Figure 4, participants in the margin-forecast-display condition significantly underestimated the chances that more marbles were red than green in all four scenarios (all $ps < .001$), with an average underestimate of -19.15 percentage points ($M_{\text{estimated}} = 64.60$ vs. $M_{\text{actual}} = 83.75$, $d = 1.26$, $t(105) = 12.99$, $p < .001$). Conversely, participants in the chance-forecast-display conditions significantly overestimated the margin prediction of how many more marbles were red than green in all four scenarios (all $ps < .001$), with an average overestimate of 22.44 marbles ($M_{\text{estimated}} = 44.94$ vs. $M_{\text{actual}} = 22.50$, $d = 0.78$, $t(111) = 8.25$, $p < .001$). As in the prior studies, both forecast formats yielded biased estimates, but the estimates were biased in opposite directions.

A regression with clustered standard errors predicting participants' margin forecast estimates in the chance-forecast-display condition found significant sensitivity to the true margin forecast value ($\beta_{\text{TRUE}} = 0.90$, $\eta^2 = 0.083$, $t(446) = 12.44$, $p < .001$). Likewise, a regression predicting

participants' chance forecast estimates in the margin-forecast-display condition found significant sensitivity to the true chance forecast value ($\beta_{\text{TRUE}} = 0.31$, $\eta^2 = 0.024$, $t(422) = 3.97$, $p < .001$). We again found a significant interaction between estimation condition and true value ($\beta_{\text{FORMAT} \times \text{TRUE}} = 0.62$, partial $\eta^2 = 0.017$, $t(868) = 6.07$, $p < .001$) in a follow-up regression predicting estimation error. These results suggest that while participant estimates were sensitive to the differences in forecasts across sample results, there was a robust forecast format bias that was larger the more extreme the actual forecasts.

Figure 4: Estimation results based on different forecast formats in Study 5

Black diamonds represent the actual forecasts for each number of red marbles in the sample. Colored dots indicate the mean estimates in one forecast format made by participants who had seen the other forecast format, with an error bar representing the confidence interval. The colored plot shapes display the distribution of estimates. Estimated chance forecasts are shown in orange and estimated margin forecasts are shown in blue.

These results confirm that complete information, which would enable a true Bayesian to make completely accurate judgments, does not eliminate people's incompatible estimates based on chance vs. margin predictions. This suggests that the prior results involving election and sports forecasts are not explained by differing beliefs about the forecast distribution between forecasters and lay people, whose beliefs arguably could have been more accurate or accurate conditional on the information available to them. Furthermore, in contrast with the previous studies, the forecasts in Study 5 were based on small samples (20 marbles) and should have been less subject to a "law of small numbers" bias (Tversky & Kahnemann, 1971) than the prior studies. Nevertheless, participants with full information overestimated the margins when shown chance forecasts and underestimated the chances when shown the margins.

As in the prior studies, the chance-margin discrepancy was pervasive across participants—63% of participants in the chance-forecast-display condition overestimated the margins of red vs. green marbles, averaging across the four tasks, while 88% of participants in the margin-forecast-display condition underestimated the chances of more red than green marbles, and was not explained by participants giving the same estimate as the forecast they had been shown.

General Discussion

Our research demonstrates that equivalent forecasts are seen as more extreme when presented in terms of the chances of an outcome (e.g., the probability of a candidate winning an election) rather than in terms of the predicted margins (e.g., the expected vote share of each candidate). The findings, from five experiments plus nine replication experiments, are difficult to reconcile with the view of

people as skilled intuitive statisticians who are adept at Bayesian reasoning, but are consistent with prior research on the difficulty of understanding statistics and non-linear relationships.

In Bayesian terms, the findings suggest a systematic discrepancy between beliefs about the expected value and tail-mass of the posterior distribution, particularly when considering the more extreme ends of the tails. The flaws in Bayesian reasoning about chances and margins identified in this research may reflect an active reliance on a miss-specified subjective distribution (e.g., with overestimated tails; Moore, Carter & Yang 2015; Urminsky 2014) on which people's chance and margin beliefs are, in fact, consistent. Alternatively, people may simply use different heuristics when interpreting chances and margins. It would be useful for future research to investigate the cognitive processes underlying the systematic discrepancy we have documented.

These findings suggest that the supposedly irrelevant choice of format for a forecast can actually have a meaningful impact on people's interpretations of the forecast and can even shift attitudes and behaviors. Given that both forecast formats consistently yielded biased estimates, but in opposite directions, it is not clear that either format is better or should be favored in communicating information. Instead, our results suggest that the most prudent practice may be to communicate both formats of forecasts simultaneously. Doing so may reduce the potential for bias, assuming that people attend to both sources of information and incorporate both in their subjective beliefs, rather than primarily focusing on and evaluating only one of the cues (Shen & Urminsky, 2013).

The implications for readers of forecasts are similar to other cognitive biases. Readers should be aware that chance forecasts are likely to seem particularly extreme, while margin forecasts may seem to represent small differences. Readers can then actively adjust their potentially biased beliefs in the reverse direction. Readers may recognize the potential for bias more readily when both

forecasts are presented together, triggering conflicting intuitions that the reader must reconcile. While forecasting sites, including fivethirtyeight.com (Figure 5), tend to do this already, news reports often present only one format or the other.

Figure 5: Communicating forecasts in both formats (adopted from: fivethirtyeight.com)

When only one format of forecast is widely available, our findings suggest that a systematic bias may result. Could this bias affect election results? For some participants, the effect of forecast format on attitudes plausibly could affect their intention to vote. However, a presidential election is a high-profile event involving substantial news coverage and personal conversations. As a result, many people are likely to have already formed behavioral intentions about whether or not they will vote and about any other election-related actions they will take, so they should have limited sensitivity to manipulated cues (Nickerson & Rogers, 2010), such as forecast framing. Furthermore, people choose to vote for a variety of reasons beyond the very low probability that their vote will change the election outcome (see Harder & Krosnick, 2008 for a review).

Taking into accounts that we found the weakest effects of forecast format when the election was closer (in Study 2), forecast format is unlikely to have a causal effect on election outcomes. In

Study 1 and a second election study in which people are presented with information about changes in chance or margin forecasts over time (see Appendix) as well as in waves 4 and 5 in Study 2, we tested the effects of forecast format on intended election behaviors. Forecast format yielded only a non-significant difference in the self-reported likelihood of voting (Study 1: $M_{\text{chance-forecast-display}} = 4.20$ vs. $M_{\text{margin-forecast-display}} = 4.46$, $t(223) = 1.47$, $d = -0.20$, $p = .143$) and in intentions for other behaviors in both studies (see Appendix). While it would be reasonable to only observe an effect of format in states where the state-level election was close and respondents' votes would be more likely to be pivotal, we do not find an effect moderated by state.

These results do not rule out the possibility of a small effect of forecast format on voting intent that our sample size was not large enough to detect. However, we expect that downstream effects on behavior are more likely when people have not already formed behavioral intentions. This is consistent with the effect of forecast format on novel behavior (an unexpected opportunity to bet on an NBA game) that we found in Study 4 (and see additional tests in the Appendix).

In contrast, we might expect to see forecast format play a meaningful role in policy and media discussions of gerrymandering (creating districts to optimize one party's electoral chances). These discussions tend not to involve quantified probabilities of a party winning a gerrymandered district. Instead, both media discussions (Ingraham, 2015) and proposed tests (McGhee, 2014) of gerrymandering tend to focus on the margin-like (and more easily quantified) issue of how the district boundaries impact relative vote shares. When newspaper readers learn that a gerrymandered district has a 60%-40% split between voters of two political parties, for example, they may see the district as more competitive than it is, underestimating the chances that the dominant party will win the election. As a result, readers may erroneously conclude that gerrymandering has less of an impact on election outcomes than it actually does, eroding support for policies intended to curb

gerrymandering. As this example illustrates, our perception of the future may depend on the format of forecasts as much as on the forecasts themselves.

References

- Bowers, J. S., & Davis, C. J. (2012). Bayesian just-so stories in psychology and neuroscience. *Psychological Bulletin*, 138(3), 389–414.
- Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science*, 17(9), 767–773.
- Griffiths, T. L., Chater, N., Norris, D., & Pouget, A. (2012). How the Bayesians got their beliefs (and what those beliefs actually are): Comment on Bowers and Davis (2012). *Psychological Bulletin*, 138(3), 415–422.
- Harder, J., & Krosnick, J. A. (2008). Why do people vote? A psychological analysis of the causes of voter turnout. *Journal of Social Issues*, 64(3), 525–549.
- Hook, J. (2016, July 17). Hillary Clinton maintains 5-point lead over Donald Trump. *The Wall Street Journal*. Retrieved from <https://www.wsj.com>
- Ingraham, C. (2015, March 1). This is the best explanation of gerrymandering you will ever see. *The Washington Post*. Retrieved from <https://www.washingtonpost.com>
- Jones, M., & Love, B. C. (2011). Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behavioral and Brain Sciences*, 34(4), 169–231.
- Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3(3), 430–454.
- Katz, J. (2016, July 19). Introducing the Upshot’s presidential prediction model. *The New York Times*. Retrieved from <https://www.nytimes.com>
- Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *TRENDS in Neurosciences*, 27(12), 712–719.

- Kühberger, A. (1998). The influence of framing on risky decisions: A meta-analysis. *Organizational Behavior and Human Decision Processes*, 75(1), 23–55.
- Larrick, R. P., & Soll, J. B. (2008). The MPG illusion. *Science*, 320(5883), 1593–1594.
- Levin, I. P., & Gaeth, G. J. (1988). How consumers are affected by the framing of attribute information before and after consuming the product. *Journal of Consumer Research*, 15(3), 374–378.
- Lieder, F., Griffiths, T., & Goodman, N. (2012). Burn-in, bias, and the rationality of anchoring. In F. Pereira, C. J. C. Burges, L. Bottou & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems 25* (pp. 2690–2798). Red Hook, NY: Curran Associates.
- Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, 9(11), 1432–1438.
- McGhee, E. (2014). Measuring partisan bias in single-member district electoral systems. *Legislative Studies Quarterly*, 39(1), 55–85.
- Moore, D. A., Carter, A. B., & Yang, H. H. (2015). Wide of the mark: Evidence on the underlying causes of overprecision in judgment. *Organizational Behavior and Human Decision Processes*, 131, 110–120.
- Nickerson, D. W., & Rogers, T. (2010). Do you have a voting plan? Implementation intentions, voter turnout, and organic plan making. *Psychological Science*, 21(2), 194–199.
- Peterson, C. R., DuCharme, W. M., & Edwards, W. (1968). Sampling distributions and probability revisions. *Journal of Experimental Psychology*, 76(2), 236–243.
- Pouget, A., Beck, J. M., Ma, W. J., & Latham, P. E. (2013). Probabilistic brains: Knowns and unknowns. *Nature Neuroscience*, 16(9), 1170–1178.

- Shen, L., & Urminsky, O. (2013). Making sense of nonsense: The visual salience of units determines sensitivity to magnitude. *Psychological Science*, 24(3), 297–304.
- Spann, M., & Skiera, B. (2009). Sports forecasting: A comparison of the forecast accuracy of prediction markets, betting odds and tipsters. *Journal of Forecasting*, 28(1), 55–72.
- Tversky, A., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin*, 76(2), 105–110.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481), 453–458.
- Urminsky, O. (2014). Misestimating probability distributions of repeated events. *Cognitive Science Society Proceedings 2014*
- Wheeler, G., & Beach, L. B. (1968). Subjective sampling distributions and conservatism. *Organizational Behavior and Human Performance*, 3(1), 36–46.
- Wright, F. A., & Wright, A. A. (2018). How surprising was Trump's victory? Evaluations of the 2016 U.S. presidential election and a new poll aggregation model. *Electoral Studies*, 54, 81–89.

Contributions

OU and LS developed the study concept and contributed to the study design. OU carried out the data collection and conducted the data analyses. LS and OU worked on the manuscript together. Both authors approved the final version of the manuscript for submission.

Competing interests

The authors declare no competing interests.

Data availability

The data that supports the findings of this study are available on OSF at <https://osf.io/f4ys6/>.

Code availability

The R code for the data analysis that supports the findings of this study are available on OSF at <https://osf.io/f4ys6/>.

APPENDIX

We briefly report the results of relevant additional measures in the paper's studies (S1-S5) and summarize nine additional studies (A1 – A9), representing all relevant data not reported in the paper. Full data, code and full stimuli are posted on OSF, at <https://osf.io/f4ys6/>.

1. Impact of forecast format on election attitudes and intentions.

Study A1 (N = 198 valid completes) was conducted on November 2nd, six days before the presidential election. AMT participants read about changes over the last 10 days in the election forecast, framed as either a shift in the chances (e.g., from an 85% chance of Clinton winning to 69%) or shift in the margins (e.g., from 53% vote share for Clinton to 52%). Participants rated this change on the same scale as in Study 2 (1 = “*very good news*” to 7 = “*very bad news*”) and were asked several questions about election-related behavioral intentions.

As in Study 1, participants who were shown the change in chance forecasts gave more extreme assessments of the news than those shown the change in margin forecasts ($M_s = 1.90$ vs. 1.36 , $d = 0.46$, $t(196) = 3.03$, $p = .001$). However, the forecast format did not yield a significant difference in self-reported likelihood of voting ($M_{\text{chance-display}} = 4.44$ vs. $M_{\text{margin-display}} = 4.39$, $t(196) = 0.70$, $d = 0.044$, $p = .758$). Tests of the differences between forecast format conditions for all behavioral intention measures in Studies 1 and A1, as well as in Study 2d (Sept. 20, N = 198) and Study 2e (Oct. 24, N = 233) are shown in Table 1 below. We also coded how close the election outcome was in each state, in order to test whether forecast format had an impact on behavioral intentions only in states where the election was close and the respondents were more likely to assume that their vote was pivotal. We found no evidence of this predicted interaction for any of the measures.

2. Forecast format and across-state election forecasts.

Study A2 (N = 155 valid completes) parallels Study 2, but we showed participants forecasts for four different non-identified U.S. states, with varying actual support for the candidates. Participants' margin estimates for the leading candidate in the chance-forecast-display condition were predicted by the actual forecasts across states (average $\beta = 0.28$, $p < .001$) but were overestimated ($\Delta = +7.02$, $d = 0.89$, $p < .001$). Participants' chance estimates for the leading candidate in the margin-forecast-display condition were predicted by the actual forecasts (average

$\beta = 3.93$, $p < .001$) but were underestimated ($\Delta = -13.83$, $d = 2.25$, $p < .001$). These finding replicated in all four states for both chance estimates ($ps < .001$) and margin estimates ($ps < .034$).

Table 1: Effect of forecast format on average election attitudes and behavioral intentions

Measure	Study	Saw Chance	Saw Margin	<i>d</i>	<i>P</i>
Election attitude*	S1	1.76	1.38	.321	.017
	A2	1.90	1.36	.461	.001
	<i>Meta**</i>	<i>1.83</i>	<i>1.37</i>	<i>.388</i>	<i><.001</i>
Optimistic*	S2d	1.16	1.15	.022	.878
	S2e	1.23	1.10	.180	.174
	<i>Meta**</i>	<i>1.20</i>	<i>1.12</i>	<i>.104</i>	<i>.282</i>
Worried*	S2d	1.15	1.11	.064	.651
	S2e	1.30	1.19	.142	.280
	S1	1.12	1.11	.010	.941
	A2	1.14	1.09	.063	.660
	<i>Meta**</i>	<i>1.18</i>	<i>1.13</i>	<i>.071</i>	<i>.840</i>
Likely to vote	S2d	3.92	4.21	-.207	.146
	S2e	4.16	4.20	-.033	.799
	S1	4.20	4.46	-.196	.143
	A2	4.44	4.39	.044	.758
	<i>Meta**</i>	<i>4.18</i>	<i>4.31</i>	<i>-.100</i>	<i>.136</i>
Donate/volunteer	S2d	1.45	1.63	-.173	.223
	S2e	1.66	1.41	.248	.059
	<i>Meta**</i>	<i>1.57</i>	<i>1.51</i>	<i>.055</i>	<i>.570</i>
Remind friend to vote	S1	3.43	3.52	-.055	.679
	A2	3.77	3.16	.385	.007
	<i>Meta**</i>	<i>3.59</i>	<i>3.35</i>	<i>.154</i>	<i>.115</i>
Post online about candidate	S2	2.18	2.49	-.204	.128
Talk to others about candidate	S2	2.57	2.78	-.139	.299

*coded as difference from midpoint of scale

**Meta-analysis results, p-values are from regression controlling for study-level fixed effects

3. Effect of forecast format on interpreting NFL forecasts.

In Study A3 (run concurrently with S2c), participants ($N = 223$ valid completes) were shown chance forecasts (71% chance of Eagles winning over the Browns) or margin forecasts (Point-spread: Eagles to win by 6.5 points) for a football game. Participants in the chance-forecast-display condition somewhat overestimated the point-spread (a non-significant difference in the predicted direction: $\Delta = +0.22$, $p = .32$), but those in the margin-forecast-display condition significantly underestimated the chances of winning ($\Delta = -8.81$, $p < .001$).

In Study A4, participants from the home states of the 2017 SuperBowl teams ($N = 393$ completes) were shown either chance forecasts (61% Patriots vs. Falcons) or margin forecasts (Patriots to win by 3.5 points). Participants in the chance-forecast-display condition significantly overestimated the point-spread ($\Delta = +3.61$, $p < .001$), but those in the margin-forecast-display condition somewhat underestimated the forecasted chances of winning (results are in the predicted direction but the difference is not statistically significant: $\Delta = -0.94$, $p = .397$). These results are generally parallel to our findings in Study 3.

4. Effect of forecast format on betting – Summary of all results.

In addition to Study 4, we conducted seven other studies (total $N = 1,999$) in which we tested whether sports forecast formats impacted betting decisions, as shown in Table 2. Table 2 shows the effect of forecast format on (a) how much out of \$5, at 3:1 odds (or out of \$4 with 5:1 odds in Study A4), people were willing to bet on the team forecasted to win and (b) how much more they bet on the winning than losing team. In all studies, participants made a betting decision about a single game, except in Study A9, where they allocated the \$5 across bets in two games. All studies were about NBA basketball games, except for Study A4.

Across all the studies in which people bet on a single game (i.e., omitting A9), they bet more on the winning team (as the proportion of funds available to bet) when shown a chance forecast than a margin (i.e., point-spread) forecast ($M_{\text{chance-forecast-display}} = .49$ vs. $M_{\text{margin forecast-display}} = .41$, $t(2118) = 4.50$, $d = 0.19$, $p < .001$). Furthermore, the effect of forecast format was stronger for games in which the forecasted chances were more extreme (interaction $\beta = 0.003$, $t(2116) = 2.32$, $p = .020$). We also find this moderation result as a marginally significant pre-registered interaction test in Study A8. We find no moderation by age, gender, knowledge of the sport or intentions to watch the game in question (all $ps > .250$).

Table 2: Impact of forecast format on betting decisions.

Study	Game (forecasted winner first)	N	p _{win}	M _{chance}	M _{margin}	D	p	Pre-reg
S4	Raptors vs. Hawks 3/6/18							
NBA	Bet on winner	371	.93	\$3.03	\$2.62	.225	.033	8906
A4	Patriots vs. Falcons, 2/5/ 2017							
NFL	Bet on winner	393	.61	\$2.16	\$2.17	-.005	.963	
	Difference in bets			+\$2.06	+\$2.03	.016	.877	
A5	Raptors vs. Hawks 3/6/18							
NBA	Bet on winner	225	.93	\$2.96	\$2.71	.124	.353	8892
A6	Trail Blazers vs. Kings, 2/27/18							
NBA	Bet on winner	502	.90	\$1.93	\$1.44	.265	.003	
	Difference in bets			+\$0.39	+\$0.10	.094	.294	8739
A7	Raptors vs. Pistons, 2/26/18							
NBA	Bet on winner	183	.91	\$2.23	\$1.28	.537	<.001	
	Difference in bets			+\$0.98	\$0	.347	.020	8683
A8	Thunder vs. Bulls, 11/15/17							
NBA	Bet on winner	217	.91	\$2.85	\$2.27	.301	.028	
	Difference in bets			+\$2.36	+\$1.70	.285	.037	
	Heat vs. Wizards, 11/15/17							
	Bet on winner	229	.56	\$1.50	\$1.36	.096	.470	
	Difference in bets			+\$0.06	+\$0.17	-.049	.715	
	<i>Difference in difference</i>			+\$2.30	+\$1.53	.172	.072	6762
A9	Thunder vs. Bulls, 11/15/17							
NBA	Bet on winner	204	.91	\$3.40	\$2.98	.226	.108	
	Difference in bets			+\$2.89	+\$2.38	.186	.186	
	Heat vs. Wizards, 11/15/17							
	Bet on winner		.56	\$0.61	\$0.65	-.043	.757	
	Difference in bets			+\$0.13	-\$0.12	.171	.224	
	<i>Difference in difference</i>			+\$2.76	+\$2.50	.086	.540	6778

The pre-reg column includes the AsPredicted pre-registration number and specifies the statistical test that was pre-registered.

5. Derivation of chance and margin forecasts in Study 5.

In Study 5, we designed the scenarios such that a Bayesian would be able to compute the correct forecast from the information provided, as follows. Let $n \sim U[0, N]$ be the unknown number of red marbles in a jar ($N-n$ marbles are green, for a total of N). A sample of S marbles are drawn with replacement, finding some number s of red marbles in the sample. The expected number of red marbles in the jar is then the expected value of the posterior distribution, conditional on the sample outcome s . The probability of a majority of marbles in the jar being red, is the mass of the posterior distribution at or above $N/2$.

Defining i as a given outcome (the number of red marbles in the sample), we want to compute $E(n | s = i)$ and $P(n > N/2 | s = i)$. We use Bayes' Rule:

$$P(n = j | s = i) = \frac{P(s = i | n = j)P(n = j)}{P(s = i)}$$

Since n is drawn from a uniform distribution, for all integer values of j in $[0, N]$:

$$P(n = j) = \frac{1}{N+1}$$

By the definition of the binomial distribution:

$$P(s = i | n = j) = \binom{S}{i} \left(\frac{j}{N}\right)^i \left(1 - \frac{j}{N}\right)^{S-i}$$

Lastly, we can sum over these terms to compute the denominator:

$$P(s = i) = \sum_{k=0}^N P(s = i | n = k)P(n = k)$$

The expected number of red marbles for the margin forecast is therefore given by:

$$E(n | s = i) = \sum_{r=0}^N r * P(n = r | s = i)$$

The probability that at least half of the marbles are red, for the chance forecast, is:

$$P(n \geq \frac{N}{2} | s = i) = \sum_{r=N/2}^N P(n = r | s = i)$$

We computed the relevant quantities for $N = 99$ and $S = 20$ in R. The stimuli shown to participants were based on sample outcomes of $i = 11$ through $i = 14$ out of 20 marbles being red.

Table 3: Statistical inferences based on samples.

Number of red marbles i in sample, out of 20	11	12	13	14
Expected number of red marbles, $E(n s = i)$	54.0	58.5	63.0	67.5
Expected margin	9	18	27	36
Probability that most marbles are red, $P(n > N/2 s = i)$	.6682	.8084	.9055	.9609

6. Additional study details

- Participants were assigned to either S2b, A1, or another study which measured estimates in the absence of any forecasts.
- Participants were assigned to either S2c, A3, or a marble-scenario. Participants were assigned to either S2d, A3, or a marble-scenario. The marble scenarios yielded similar results but were not included because they contrasted sample observations with chance forecasts (as opposed to margin forecasts vs. chance forecasts, as in S5).
- Participants were assigned to either S2f, or another study which tested the effects of anchoring participants on numbers equivalent to the forecasts or the actual values, but not identified as such.

7. Open science materials

The pre-registrations associated with the studies can be found online as follows:

Study 3 (AsPredicted #6391): <http://aspredicted.org/blind.php?x=ba38f8>

Study 4 (AsPredicted #8906): <http://aspredicted.org/blind.php?x=5qw89v>

Study 5 (AsPredicted #6686): <http://aspredicted.org/blind.php?x=gq8hm3>

Study A5 (AsPredicted #8892): <http://aspredicted.org/blind.php?x=y9pd7c>

Study A6 (AsPredicted #8739): <http://aspredicted.org/blind.php?x=jx9kt7>

Study A7 (AsPredicted #8683): <http://aspredicted.org/blind.php?x=7sv8xw>

Study A8 (AsPredicted #6762): <http://aspredicted.org/blind.php?x=w8p4z9>

Study A9 (AsPredicted #6778): <http://aspredicted.org/blind.php?x=qp7xn7>

In addition, data files, survey materials, pre-registrations and analysis code (in R) for the studies are available on OSF (<https://osf.io/f4ys6/>).