

Experimental Electionomics: How Election Forecasts Influence Voter Turnout

Ben Kaplan
Carnegie Mellon University

Advisor: Dr. Geoffrey McGovern

Abstract

We create forecasts as a way of reducing the uncertainty that surrounds uncertain events. Often, forecasts provide information that can be acted upon, and in the case of elections, their use as a decision aid can have a direct effect on the outcome of the election. Currently, it's unclear how potential voters incorporate election forecasts in their decision making. To gain a better understanding of this, a survey was created and distributed to 186 subjects. Participants were introduced to two fictional candidates in a hypothetical presidential election, and provided their stances on a series of political issues. Subjects were asked whether or not they'd vote for their preferred candidate, and were randomly shown one of four different election forecasts. Then, they were asked again whether or not they'd vote for their preferred candidate. Overall, participants' decision to turn out switched 14% of the time after being shown an election forecast. There was no evidence to suggest that there is a relationship between the type of forecast shown and the likelihood that an individual switched his or her turnout decision. This research shows that election forecasts *do* have an impact voting behavior, but the mechanism by which this occurs remains unclear. As such, future research on this topic should be conducted to better understand how voters use forecasts. This will assist media outlets and campaigns in being well informed of the impact associated with forecast reporting, allowing American democracy to grow stronger.

Table of Contents

1. Introduction	3
2. Background	5
2.1 Election Forecasting	6
2.2 Voter Turnout	8
2.3 Race Closeness and Voter Turnout	10
2.4 Election Forecasts as Decision Aids	10
2.5 Significance	12
3. Scenarios and Hypotheses	15
3.1 Close Races	15
3.2 Landslide Races	16
3.3 Hypothesis 1: Close Races	18
3.4 Hypothesis 2: Landslide Races	18
4. Methodology	19
5. Results	24
6. Discussion	27
6.1 Explanation 1: Numeracy	27
6.2 Explanation 2: Survey Bias	28
7. Implications	31
8. Limitations and Future Work	34
9. Conclusion	38
10. References	39

1. Introduction

We create forecasts as a way of reducing the uncertainty that surrounds uncertain events. When it comes to uncertain events like the weather, sports, or even elections, humanity is impatient. We want to know what will happen and, naturally, since we can't know about an outcome until it occurs, we sate our hunger for certainty through the next best thing: a forecast. Since we don't have access to time machines, it's incredibly difficult to know what to expect in the future, and even more difficult to anticipate something with a high level of certainty. Forecasts allow us to take pieces of information related to an uncertain event and piece them together in a way that helps us gain clearer insight into reality. In most contexts, there's a small set of possible outcomes, but there are many more pressures and factors that actually have an effect on those outcomes. It's for this reason that forecasting both accurately, as well as consistently, is no easy task. This is especially the case for election forecasting. Unlike other forecasting contexts, elections don't take place in a vacuum. The results of an election are directly affected by the action – or inaction – of thousands, and sometimes millions, of individuals.

The results of an election are a function of two main factors: the number of people that support each candidate, and the number of people that actually vote in the election. Traditional election forecasts tend to focus on determining the former, but not the latter. Good forecasts attempt to account for both. It's much more difficult to know for sure whether or not an individual will vote than it is to know which candidate an individual supports. When forecasts are inaccurate, and systematic bias doesn't appear to be the culprit, a likely reason for the

inaccuracy is that the forecast doesn't properly account for voter turnout. But let's cut forecasters some slack. Turnout is incredibly complicated, and Political Science literature seems to disagree on the specific components and factors that influence it. Distance from a polling station, political affiliation, age, race, socioeconomic status, mail-in voting laws, etc. are all known to be associated with voter turnout. But the waters are murky.

The research discussed in this paper aims to shed light on the effects of a factor that I hypothesize influences voter turnout: election forecasts, particularly those for U.S. elections. Thus, at the center of my efforts, I aim to determine if individuals use election forecasts as decision aids when deciding whether to vote. In the case that they do, I wish to determine how, and under what circumstances, forecasts are used as a decision aid.

Forecasts can be used as decision aids in other contexts. For instance, when one is not beholden to any particular college basketball team during March Madness, it's reasonable to believe that they would wish to watch games that will be most exciting. If we assume that excitement can be measured in terms of the closeness of a game, a forecast can act as an effective tool to determine which game to tune in to. The difference between using forecasts to determine which March Madness games to watch, and using forecasts to determine whether or not to vote on election day, is that the results of the March Madness game will not change just because you decided to tune in. The same can't always be said for elections and their subsequent results. In an election context, if forecasts have a direct influence on the results, there may be troubling implications. Bad actors could decide to influence turnout by promoting inaccurate forecasts. This becomes a dilemma, because without understanding the model's methodology, there may be no detectable difference between an intentionally inaccurate forecast and a

naturally inaccurate one. By better understanding how potential voters use election forecasts in their decision making, campaigns can adapt, pollsters and researchers will have an ethical justification for enhancing our forecasts, and our democracy can grow stronger.

2. Background

Election polling has existed for nearly two centuries. The first documented U.S. poll was in 1824 by *The Aru Pennsylvanian* in 1824, showing Andrew Jackson leading John Quincy Adams by 335 votes to 169 in the presidential election (Madonna and Young, 2017). A century later, *Literary Digest* became America's most prominent pollster, conducting polls that sampled huge portions of the population. Even by today's standards, these polls, conducted with postcards, are some of largest polls ever. *Literary Digest* was able to correctly predict the winner of the 1920, 1924, 1928, and 1932 elections. However, due to a huge polling error that came as a result of polling 2.3 million affluent Americans, they obtained a biased (non-representative) sample. This is because they sampled only those that read *Literary Digest* and had the ability to respond to it, and these people tended to be wealthier than the entire U.S. population (Squire, 1998). Consequently, they wrongly predicted that Alf Landon would win the race by a large margin, when in fact, it was Franklin Roosevelt that won by a large margin. This is arguably the largest polling blunder of all time and it led to a massive demand of improvements to polling. Thus, it wasn't long before *Literary Digest* became irrelevant and other pollsters, such as George Gallup and Elmo Roper, who used demographically representative samples, were able to successfully predict Franklin Roosevelt's 1936 victory and became reliable pollsters for decades to come (Gosnell, 1937).

2.1 Election Forecasting

Uncertainty is a great cause of stress and fear. In an attempt to reduce and ultimately eliminate this great cause of stress, we've developed forecasting methods that allow us to become more certain about uncertain events. Over time, forecasting methods improve, making forecasts more accurate. However, this does not mean that they're always correct. Research shows that forecasts have become more accurate over time, but errors still do happen at a consistent rate (Wlezien and Jennings, 2017). When forecasts are wrong, this is not necessarily indicative of a poor model. Forecasts output a percentage chance of a singleton event occurring. Since these events generally only occur once (this is always true for elections), it's difficult to understand whether a forecast is accurate or not and if it is accurate, whether or not it's accurate for the right reasons.

It is a lot more difficult to understand the accuracy of a model for a presidential race, than that of a model for congressional races. This is because congressional forecasts can be evaluated across dozens of races, therefore making it easier to understand whether or not there's systematic bias, as well as what the cause of that bias might be. Presidential races are one and done. This is why it's difficult to say whether or not a model like *FiveThirtyEight's 2016 Presidential Election Forecast* was accurate in giving Donald Trump a 28.6% chance of winning the presidency on election day 2016. On the other hand, one can more safely conclude that *FiveThirtyEight's 2018 House Forecast* was accurate in predicting that the Democrats would pick up 39 seats in the House of Representatives (in reality, Democrats won 40 seats). This is because across 435 races, there was no significant bias in either direction.

In the beginning, most forecasts were constructed through polling. Over time, polling methods improved, with a focus on reducing bias by standardizing the construction of polling questions, as well as improving sampling methods to sample likely voters, as opposed to registered voters (Zaller and Feldman, 1992). More recently, efforts to reduce the costs of polling while maintaining accuracy have been a large research focus, with methods such as multilinear regression and poststratification (MRP) at the forefront. MRP is used to create accurate non-representative polls. Under certain circumstances, MRP makes internet polling a cheap and viable option to obtain accurate results, with a much larger sample size than traditional telephone (through random digit dialing) or postal polls (Goel, Obeng, & Rothschild, 2015; Wang et al., 2015; Konitzer et al., 2017).

A singular poll often does not make for an accurate forecast. To improve forecast accuracy, it's typical for polls to be aggregated. To do this, poll aggregators apply weights to polls based on recency, sample size, sampling type, pollster credibility, and the pollster's historical bias, and aggregate them to create an overall forecast of a race. More complex models use factors along with polling aggregates, such as fundraising, incumbency advantage, historical trends, demographic trends, turnout predictions, expert opinions, and other predictors.

Under the appropriate circumstances, social media sites are one of the best ways to gauge public opinion. This is because people are more likely to say what they feel, freely have discourse, and perform actions that signal what their political leanings or interests may be, such as liking or commenting on posts or following people or pages. By using machine learning algorithms, researchers have been able to predict factors such as candidate support, favorability, and voter turnout, and ultimately build forecasts based on these predictors. Natural language

processing (NLP) techniques such as topic identification and sentiment analysis are used to interpret data from social media to understand how individuals feel about particular candidates (Marchetti-Bowick and Chambers, 2012). Simpler approaches involve taking the overall number of tweets that mention a particular party or candidate during a particular time-frame (the week before the election, for example) and using the percentage of each candidate's/party's mentions as a way to predict voter share. In the limited cases where this method is used, it's shown to be promising, with a mean absolute error of 1.65% (Tumasjan et al., 2011). While these methods are in the fairly early stages of research and development, they appear to have potential moving forward. The benefits of these methods are that they're relatively cheap and they're more accurate than traditional polling months and weeks away from an election.

2.2 Voter Turnout

To understand whether or not distance affects voter turnout, a clear understanding of the leading normative and descriptive models of voter turnout is necessary. The normative model is known as rational choice theory and the descriptive model takes a modified cost-benefit approach. Rational choice theory was first described in 1957 by Anthony Downs, who argues that it is irrational to vote in first-past-the-post elections. The logic of his argument is as follows: assuming that every voter is a rational decision maker, the probability that their vote will be the pivotal vote (meaning that that vote singlehandedly decides the outcome of the election) in a first-past-the-post election is incredibly low (Downs, 1957; Aldrich, 1993; Riker, 1995). Downs claims that since there is a time cost associated with the act of voting, it is irrational to accept this

cost and vote in an election, knowing that it is incredibly unlikely that the vote will have an impact on the election outcome.

Rational choice theory is effective in that it lays out an argument for why a voter should use a type of cost-benefit analysis relating to voting, but fails to describe how we actually decide whether or not to vote. As we know, elections have thousands or even millions of voters, meaning rational choice theory does not adequately explain how people decide whether or not to vote. Instead, a modified cost-benefit approach helps to paint a clearer picture of how voters might actually decide whether or not to vote. This theory, proposed in the years following rational choice theory, argues that an individual will decide to participate in an election so long as the costs of voting are lower than the benefits of voting. This can be modelled by the following equation (Haspel, 2005):

$$C < P * B + D$$

where C refers to the costs of voting, P refers to the probability of placing the pivotal vote, B represents the benefits that arise from the favored candidate winning, D represents the satisfaction of performing one's civic duty (Riker and Ordeshook, 1968; Ferejohn and Fiorina, 1974; Niemi, 1976; Sanders, 1980). Most of the existing literature is focused on understanding the benefits, costs, and civic duty terms in this equation, but existing research often overlooks the P term. Particularly, how the perception of the probability of a pivotal vote plays a role in turnout decisions. Uniquely, election forecasts can impact how we perceive this probability, and as such, it's expected that an election forecast can have an impact on voter turnout. Until now, however, there isn't really any literature that appears to investigate whether or not this is true, as well as the mechanisms and scenarios where forecasts matter to potential voters.

2.3 Race Closeness and Voter Turnout

While there is a limited amount of existing research on the impact of election forecasts on voter turnout, there is literature that explores the relationship between race closeness and voter turnout. There's clear similarities between these research topics. In the case of race closeness research, there's no assurance that an individual's perception of closeness is aligned with the reality of the state of the race. Election forecasts help to bridge that potential gap between perception and reality, as they're able to give a clearer insight into the state of the race. The implication of this may be that race closeness, in the context of election forecasts, has a stronger effect on voter turnout than race closeness does. Existing research tends to disagree about the impact of race closeness on voter turnout, with some research stating that there's no statistically significant effect of race closeness on voter turnout, and other research stating that it has been shown to have a small effect (Cox et al., 1989; Matsusaka, 1993; Endersby et al., 2002). One such study actually investigates how forecasts shown during election-night television broadcasts impact voter turnout in a close presidential race, and did not find any evidence that these forecasts impacted turnout (Tuchman and Coffin, 1971). Like research that focuses on election forecasts and voter turnout, the depth of research into this topic is limited, and there's a lot more to uncover and learn.

2.4 Election Forecasts as Decision Aids

After the 2016 presidential election, researchers have been interested in understanding how potential voters are influenced by election forecasts. Research in this topic aims to obtain

answers to a few questions. First, how well do potential voters actually understand the forecasts that they see? Second, how do voters incorporate these forecasts in their decision process when deciding to vote? And third, what types of situations are electoral environments lead to increases or decreases in overall voter turnout?

Studies have been performed to better understand numeracy within the context of election forecasts. One of the main findings is that people tend to struggle with distinguishing between forecasts that predict the vote share and forecasts that predict a candidate's chances of winning a race (Westwood et al., 2019). Both of these forecast types convey the same information, but the latter presents it in a fashion that can be more easily translated to a decision. The issue is that in practice, people often don't pay attention to these details, leading to suboptimal decision making. In addition to this, people aren't as skilled at understanding probabilistic forecasts in an election context because elections don't happen very often. For example, a person's decision to take an umbrella with them when there's a 45% chance of rain occurs within the context of a type of forecasted event that happens every day. People have a lot of exposure to weather forecasts, but not to election forecasts, so interpreting probabilistic forecasts in unfamiliar contexts can lead to poor decision making at a higher rate (Kunreuther, Novemsky, and Kahneman, 2001).

Researchers have also been effective in understanding how forecasts play a role in how potential voters make voting decisions. This is largely a result of a large body of research surrounding voter turnout, and the various factors and influences that affect it (Ansolabehere and Iyengar, 1994; Delli Carpini, 1984; Mutz, 1998). In addition to turnout based effects, there's a growing amount of evidence linking election forecasts to the specific behaviors of campaigns.. Candidates that are polling well in certain states may want to focus their attention on states

where their poll numbers are weaker than expected, in an attempt to improve their chances and collect as many electoral college votes as possible (Enos and Hersh, 2015; Mutz, 1997). There's also evidence to link election forecasts leading up to an election to policy decisions made by administrations in order to hedge against expected risks that result from an impending election. For instance, it's been reported that Obama's lack of response regarding Russian interference in the 2016 election was partially a result of confidence in Hillary Clinton's ability to win, which could potentially nullify the public's support for these inquiries (Miller, Nakashima, and Entous, 2017) .

2.5 Significance

Since the very first election forecasts, the main goal of forecasters has been to build models that produce forecasts that are as accurate as possible. However, not much thought or research has been done on specifically how election forecasts impact voter turnout, and subsequent election results. Forecasts are picked up by the media and are shared with the public, which makes sense, because the purpose of a forecast is to reduce uncertainty about an uncertain event. But with this, an issue arises: what if the election forecasts themselves impact turnout?

Let's use the 2016 presidential election as a hypothetical to make this point more clear. In 2016, election day forecasts gave Donald Trump anywhere from a 30% chance of winning, to a 1% chance of winning (Singer-Vine, 2016). It's not that hard to believe that these forecasts that came from public media sources, such as *FiveThirtyEight*, *The New York Times*, *Slate*, and *The Huffington Post*, as well as prediction markets such as *PredictIt* and *PredictWise*, with millions of readers each day, had a non-zero impact on voter turnout. This research is driven by the

assumption that people don't only look at forecasts because they're interested in seeing how each candidate is doing in a particular race, but rather that people also use forecasts (both consciously and subconsciously) as a tool for determining whether or not they should vote on election day. For example, if *The New York Times* on election day, gives Hillary Clinton a 85% chance of winning, then one might feel like voting for her is a waste of time, since the vote is unlikely to matter. The logic here is similar to how rational choice theory operates. However, the difference in this case is that voters are given an indication of how each candidate will do in the race. If election forecasts are used by potential voters as a decision aid in determining whether or not to vote, they're likely to be used by individuals as a part of their modified cost-benefit analysis. I posit that when a particular candidate is forecasted to win with a high degree of certainty, an individual's perception of the probability of their vote being pivotal decreases, thereby decreasing that individual's likelihood of voting in the election.

In the context of the 2016 presidential election, this would mean that supporters of Hillary Clinton would be less likely to go to the polls on election day, given that she is projected to win with relatively high confidence. Likewise, but to a lesser extent, supporters of Donald Trump would be less likely (but more likely than Clinton supporters) to go to the polls, so that they can try to help the candidate who is unlikely to win (according to forecasts), to win.

This research is focused on answering three key questions: first, whether or not individuals use forecasts as a decision aid in determining whether or not to vote in an election. Second, how race closeness impacts an individual's voting intentions as a result of viewing an election forecast. Finally, understanding the overall impact of the effect of election forecasts, i.e. would the results of recent elections change if forecasts didn't exist?

If it is true that voter turnout can be impacted by election forecasts, an ethical question arises. Do election forecasts undermine American democracy? I would argue that they could. This is concerning because while policy positions, reputation, and qualifications are all important criteria when deciding which candidate to vote for, as well as whether or not to vote at all, the likelihood of the candidate you support winning should not have an impact on the choice of whether or not to vote. It's irrelevant information. It's as fallacious as saying, "I didn't vote because I didn't want to vote for the loser". The fact that it does, poses a problem, especially since election forecasts are more accessible to potential voters than ever before . The other issue, is that election forecasts don't really disclaim that they might be wrong or inaccurate (that's just an inherent risk), but often, we treat them as though they are reality or fact.

It's even more troubling that the understanding that election forecasts can influence voter turnout could lead to the creation of purposely inaccurate forecasts that serve to deter or encourage people to vote. This is where an ethical dilemma arises. Is it unethical to publish an inaccurate forecast, even if you didn't intend for the forecast to be inaccurate? This paper will discuss potential ways to deal with the effects of forecasts in the case that they're shown to influence voter turnout. The goal of this is to help curb the effects of bad actors from influencing the democratic process through the spread of misinformative election forecasts.

3. Scenarios and Hypotheses

For the purposes of this research, there are two election scenarios being explored. The first is a race that is forecasted as being “close”. I define a predicted close race as a race that is forecasted where the difference between both candidates’ forecasted chance of winning is less than 20% (a 60-40 split). The second case is a race that is forecasted as being a landslide. I define a predicted landslide as a race with forecasts that show that the difference between both candidates’ forecasted chance of winning is greater than 50% (an 75-25 split).¹ The definitions for a both close races and landslide races are based on what people would generally consider as close or not. Leading forecasters, such as *FiveThirtyEight* uses these same cut-offs for races that they determine to be “solid Democrat”, “solid Republican”, or “toss-ups” (Silver, 2018). For both close races and landslide races, there are 5 potential effects that a forecast might have on voter turnout. I’ll be discussing each one, potential explanations, and whether or not the effect is likely to occur.

3.1 *Close Races*

Potential result 1: No effect. This could be the result of election forecasts having no influence on an individual’s likelihood of voting. This would be expected under a subpopulation of people that vote because they feel that it’s their civic duty to do so. No effect in race outcome could also be seen if turnout increases or decreases at an equal rate for both the underdog and

¹ For the purposes of this paper, I’ve decided to look at only two scenarios, and this leaves a gap of races where the race is neither close nor a landslide. In an attempt to determine whether or not an effect exists, I’ve decided to look at the extremes, but as noted later, it may be worth investigating a wide range of race types and environments in the future.

frontrunner as a result of a forecast. This would imply that forecasts for close races neither harm nor help the frontrunner or underdog's chances of winning.

Potential result 2: Underdog turnout increases. In this case, a "rallying" effect would be expressed. An explanation for this effect would be that individuals recognize the impact their vote might have, and since victory for their supported candidate is within arm's reach, they value the costs of voting as less of a factor in their decision making, therefore they decide to vote. On an aggregate level, this effect could lead to the success of an underdog candidate.

Potential result 3: Frontrunner turnout increases. This case represents a "reality check" scenario, where supporters of the frontrunner candidate had at one point overestimated the support for their candidate and upon viewing the most recent forecast, realize that the race is much closer than they once thought. As such, some supporters that at one point had deemed their vote unnecessary now feel as though they need to vote for their supported candidate. As a result, turnout for the frontrunner increases.

Potential results 4 and 5: Underdog turnout decreases, frontrunner turnout decreases. Both of these effects are unexpected to occur for a close race. In order for turnout to decrease, the benefit of voting would have to decrease, however, in a close race, it would be expected to increase. These scenarios are possible but go against the traditional understanding of how voter turnout decisions are made. As such, it would be incredibly unlikely for an election forecast to have these effects on turnout.

3.2 *Landslide Races*

Potential result 1: No effect. This follows the same logic and explanation for why there might be no effect in the case of a close race.

Potential result 2: Bandwagoning (frontrunner turnout increases). In this scenario, individuals that view a forecast that show their preferred candidate as the frontrunner in an expected landslide race decide to vote as a way of “being on the winning side”. Individuals that are undecided on election day in an expected landslide race may decide to vote for the frontrunner if the costs of voting are low as a way of feeling as though they’re doing their civic duty. Even though they’re undecided and otherwise wouldn’t have supported any candidate, the nature of the race as a landslide gives a way for an individual to fulfill their civic duty without much fear of affecting the outcome of the race.

Potential result 3: Final stand (underdog turnout increases). In this scenario, individuals that view a forecast that show their preferred candidate as the underdog in an expected landslide may decide to vote in order to give their candidate a chance (although unlikely) at winning the race.

Potential result 4: “I’m not needed” (frontrunner turnout decreases). In this scenario, supporters of the frontrunner candidate that view a forecast on election day determine that it’s unlikely that their vote will matter because the race is expected to be a landslide. This effect would be more likely among those that derive a low sense of civic duty from voting and have high costs associated with voting.

Potential result 5: Accepting defeat (underdog turnout decreases). This scenario has individuals opting not to vote for their supported candidate if they’re expected to lose by a

landslide. These individuals may find that turning out to the polls just isn't worth their time, as it's unexpected that their candidate can win. This effect would be more likely among those that derive a low sense of civic duty from voting and have high costs associated with doing so.

3.3 *Hypothesis 1: Close Races*

In close races, I expect that frontrunner and underdog turnout will both increase when potential voters view election forecasts. While both increase, I hypothesize that underdog turnout will increase at a greater rate than frontrunner turnout, giving an advantage to the underdog candidate. The reason for this is that a close forecast gives a substantive notion that the race is within arms reach, motivating supporters of the underdog not to give up and to push their candidate to a majority.

3.4 *Hypothesis 2: Landslide Races*

In landslide races, I expect that frontrunner and underdog turnout will both decrease when potential voters view election forecasts. I hypothesize that frontrunner turnout will decrease at a higher rate than underdog turnout for landslide races. This is because in landslide races, the supporters of underdog candidates could be less willing to give up on their candidate than those in support of a frontrunner than is expected to win by a landslide. If this hypothesis were correct, it would also imply that in landslide races, underdogs have a larger chance of winning than predicted if potential voters are using election forecasts as decision aids, because supporters of the frontrunner will likely decide not to vote, while supporters of the underdog will likely decide to vote for their candidate, essentially doubling the value of each underdog vote (a

vote for the underdog that would have otherwise been cancelled out by a vote for the frontrunner).

4. Methodology

To understand the relationship between election forecasts and voter turnout, an experiment was conducted on Amazon's Mechanical Turk. Eligible subjects completed a survey regarding a hypothetical presidential election. Eligibility required that subjects were registered voters in the United States. As such, all subjects were at least 18 years old and were U.S. citizens. The survey had three components:

The first aims to understand basic demographic features of the respondent. This includes race, age group, gender, and whether they've voted in a recent election. The purpose of collecting this data is to understand whether or not any effects of election forecasting on voter turnout are associated with these basic demographic variables. The survey questions in this section used language that is widely used in the field of social science research for collecting demographic data so as not to bias responses.

The second component is meant to gauge the subject's political ideology. To do this, subjects are asked ten of seventeen questions taken from the *Pew Research Political Typology Quiz* (Pew Research Center, 2017). The ten questions were selected to understand answers to a wide range of political issues. For each question, a score from -1 or 1 is given to each subject depending on their answer choice. A score of -1 represents the more traditionally liberal view and a score of 1 represents the more traditionally conservative view. The sum of scores for the ten questions represents the subject's political ideology, with -10 being "extremely liberal", 0

being “centrist”, and 10 being “extremely conservative”. The purpose of this survey component is twofold. Firstly, it’s to determine whether or not any effect that election forecasts may have on voter turnout could be related to partisanship or ideology. Secondly, this data helps serves as a way to determine whether or not subjects were honest in their responses. This is because in the third component, subjects are asked to pick their preferred candidate from two choices, one with a liberal platform, the other with a conservative platform. If there is a misalignment between a person’s political ideology and their preferred candidate’s platform, it could suggest that a different factor is influencing a subject’s choice of candidate, which may jeopardize the validity of any findings.

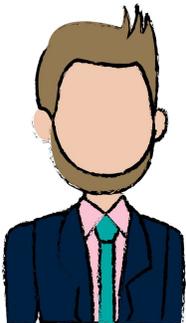
The third component is the experimental component. Subjects are given a hypothetical scenario where they’re told that it’s Election Day 2028 in the United States. Subjects are introduced to two candidates, Dan Jennings and Jim Dwyer. Subjects are shown a list of stances that each candidate has as part of their platform. One platform takes clear liberal stances while the other takes clear conservative stances. These stances come directly from the questions asked in the second component of the survey. The platforms are randomly assigned to one of the two candidates, to ensure that the names representing the candidates do not bias candidate preference in aggregate. Similarly, images of each of the candidates are shown along with their platforms (Figures 1 and 2). These images are also randomly assigned to each of the two candidates, so as not to bias subjects’ candidate preferences. Subjects are then asked to determine which of the two candidates they would support based on the information given. Next, they’re asked if they would vote for their preferred candidate. Subjects are then randomly assigned to one of four experimental conditions.

Figure 1.



One of two images randomly assigned to represent either Jim Dwyer or Dan Jennings

Figure 2.



One of two images randomly assigned to represent either Jim Dwyer or Dan Jennings

The first condition is one where their preferred candidate is the frontrunner and the race is a close race. The second condition is one where their preferred candidate is the underdog and the race is a close race. The third condition is one where their preferred candidate is the frontrunner and the race is predicted to be a landslide. The fourth condition shows their preferred

candidate as the underdog and the race is predicted to be a landslide. Depending on the condition the subject is randomly assigned to, they're shown only one of the four election forecasts (Figures 3-6). These forecasts share the design and style of *FiveThirtyEight's* 2016 presidential election forecast, to help ensure they're as easy to digest and understand as reasonably possible. Subjects are then asked "considering your typical weekly schedule, would you make time to go to the polls to vote for [preferred candidate name]?". The purpose of this component of the survey is to understand if there's a change in an individual's choice of whether or not to vote for a candidate after being introduced to an election forecast that provides insight into the status of the race.

Figure 3.



A forecast where the race is close and the preferred candidate (in this case it's Jim Dwyer) is the frontrunner.

Figure 4.



A forecast where the race is close and the preferred candidate (in this case it's Jim Dwyer) is the underdog.

Figure 5.



A forecast where the race is a predicted landslide and the preferred candidate (in this case it's Jim Dwyer) is the frontrunner.

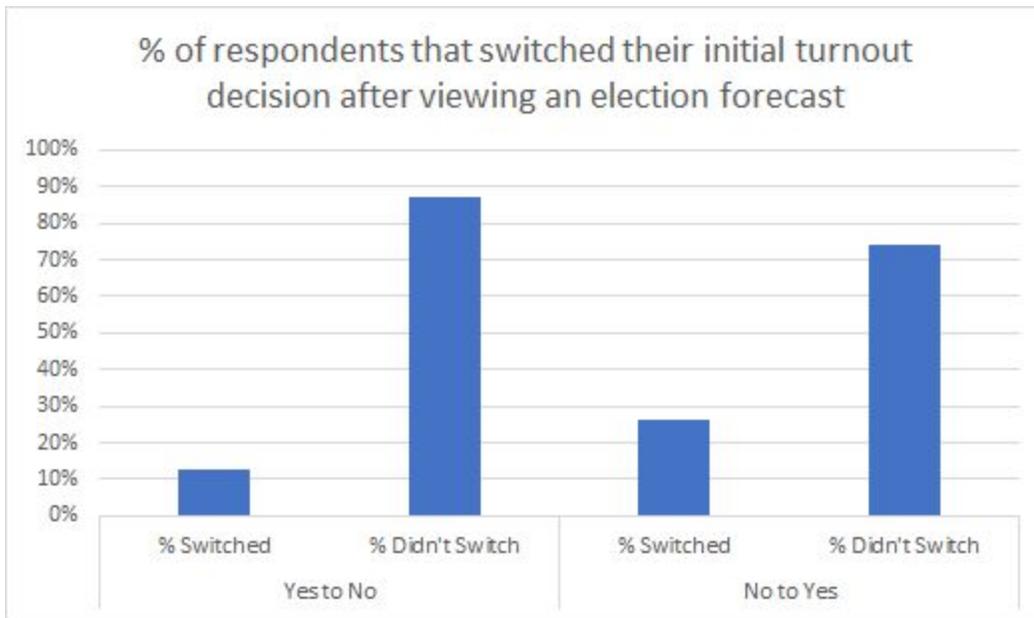
Figure 6.



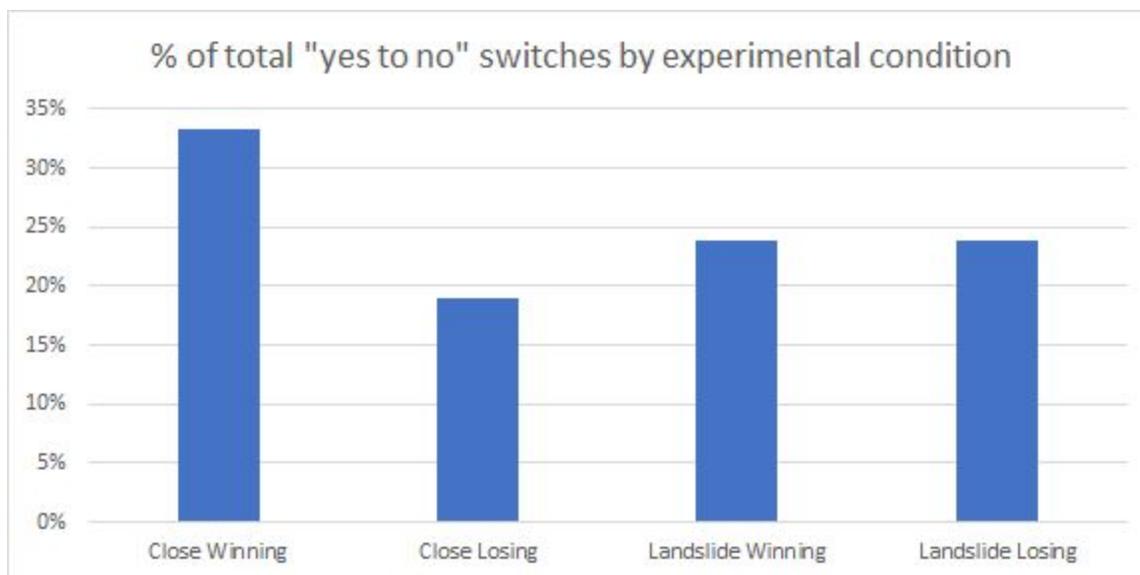
A forecast where the race is a predicted landslide and the preferred candidate (in this case it's Jim Dwyer) is the underdog.

5. Results

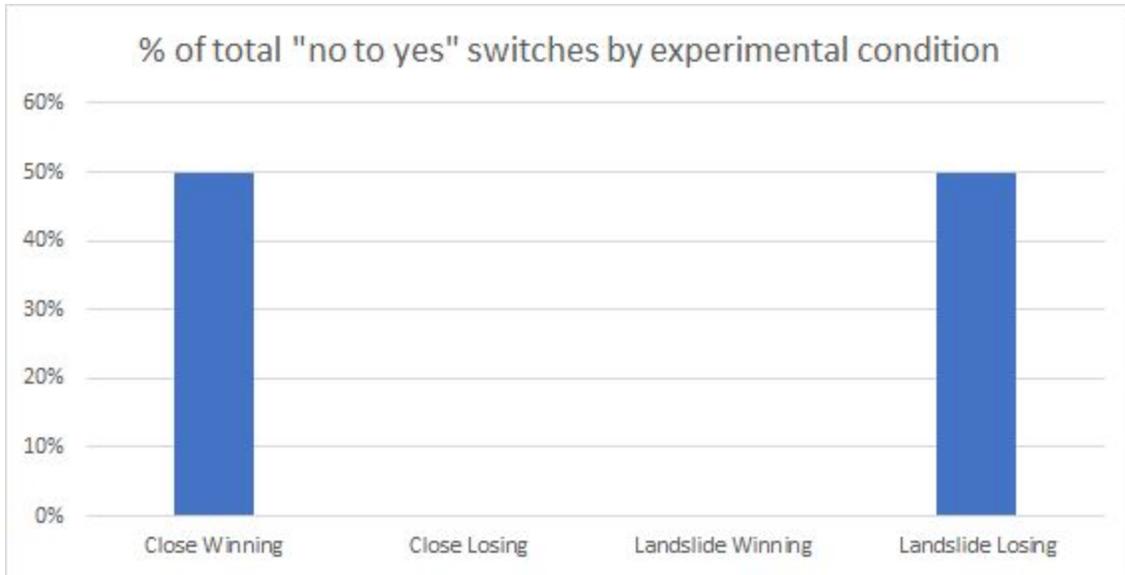
Participants (N = 186) came from Amazon's Mechanical Turk (MTurk). The majority of subjects were aged 25-34 years old. 65% of subjects identified as male. 49% of subjects were registered democrats, 27% were registered republicans, and 21% were registered independents. 83% of respondents voted in the 2018 midterm elections. 14.5% of subjects changed their initial turnout decision after being exposed to an election forecast. Turnout decision switching appeared to occur at a higher rate among those that had initially opted not to vote for their initial candidate, than those that decided to. For those that decided against voting for their preferred candidate initially (n = 23), 26% changed their decision to vote after being exposed to an election forecast. For those that decided to vote for their preferred candidate initially (n = 163), 13% opted not to vote for them after being exposed to an election forecast.



In subjects where “yes to no” turnout decision switching occurred ($n = 27$), 33% of switching occurred in races where the race was forecasted as close and the subjects’ preferred candidate was winning. 19% of switching occurred in races where the race was forecasted as close and the subjects’ preferred candidate was losing. 24% of switching occurred in races where the race was forecasted as a landslide and the subjects’ preferred candidate was winning, and 24% of “yes to no” switches occurred in races that were forecasted as a landslide with the subjects’ preferred candidate predicted to lose.



Among subjects where “no to yes” turnout decision switching occurred ($n = 6$), 50% of these switches occurred in races that were forecasted as close, where the subjects’ preferred candidate was predicted to win. The other 50% of no to yes switches occurred in races that were forecasted as a landslide, where the subjects’ preferred candidate was predicted to lose. Zero “no to yes” switches occurred in either of the other two experimental conditions.



6. Discussion

According to the experimental results, election forecasts do have an impact on voter turnout. However, this impact is not as hypothesized. In the case of subjects that switched their decision to vote from “yes” to “no”, the type of race forecasted doesn’t appear to have any statistically significant effect on subjects’ behavior. This finding is unexpected, because while a statistically significant number of subjects decided to change their decision to vote from “yes” to “no”, the type of race forecasted appears to be independent of their “yes” to “no” switch. Across all experimental conditions, there is no evidence to suggest that any one type of forecast has an effect on an individual’s decision to vote. The survey did not ask subjects that switched their votes to explain why they did, but in hindsight, this qualitative information may have been useful in understanding the mechanism behind people switching their turnout intention. If it’s not the type of forecast that impacts an individual to switch their decision to vote from “yes” to “no, then there are two possible explanations for this effect:

6.1 *Explanation 1: Numeracy*

The decision science literature states that in general, individuals often struggle when it comes to understanding numbers (Peters et al., 2006). Percentages pose a particular problem and in the case of this experiment, the forecasts were presented as percentage chances of winning the race. Just because a forecast shows Jim Dwyer with a 20.5% chance of winning the race doesn’t mean that the subject that views this forecast will fully comprehend what this means and be able to use this information effectively in their decision making. Research also shows that individuals tend to overestimate the probability of unlikely events (Barberis, 2013; Taleb, 2007). If a

forecast shows that it's unlikely for a particular candidate to win, they may misinterpret what the given probability actually means. In other words, it's possible that subjects' failed to interpret the forecasts correctly, thereby causing them to make suboptimal decisions. Considering that "yes" to "no" switching occurred fairly equally across all experimental conditions, this explanation appears to be likely.

6.2 *Explanation 2: Survey Bias*

There's a possibility that the survey question biased participants to consider factors outside of the scope of the election forecast that they had not initially considered when making their first turnout decision. The exact wording of the two questions after being introduced to the two candidates, but before the forecast is as follows: "Based on the information given, which of these two candidates do you support?" and "Would you vote for the candidate you selected in the question above?". The wording of the question after the forecast was, "It's Election Day 2028. You wake up and see the following forecast showing the chance each candidate has to win the presidential race: [insert forecast]. Considering your typical weekly schedule, would you make time to go to the polls to vote for [insert preferred candidate's name]?".

The main concern with the wording of this question is that subjects are asked to think about their typical weekly schedule and whether or not they're able to make time to go to the polls. From a survey design perspective, this question was worded to elicit thinking about the costs of voting when making their decision. This is important, because costs play a critical role in voter turnout decisions. Without getting participants to consider the costs of voting, it would be unclear how much forecasts actually impact voter turnout decisions in a real-world environment.

At the same time, this question does get participants to think about a factor that they had not previously considered prior to being introduced to the forecast and making their initial turnout decision. For this reason, there's a possibility that participants were biased by thinking about costs in only one of the two times they were asked whether or not they'd turn out. This means that the effect seen may not have much to do with election forecasts, but instead has more to do with the costs of voting.

In an attempt to further understand what mechanism may be at play that causes turnout decision switching, demographic trends among individuals of turnout decision switching were compared to the demographic trends among those subjects that did not switch their turnout decision. For “yes to no” switchers across most demographics, including age, gender, and party affiliation, there was no statistically significant relationship found between subjects' tendency to switch their turnout decision and the demographic variable. However, there was a statistically significant relationship ($p = .0273$, $\chi^2 = 4.874$) between subjects' tendency to switch their turnout decision, and whether or not they voted in the 2018 midterm elections. Subjects that voted in the 2018 midterm elections switched their turnout decision only 10.3% of the time, while subjects that did not vote in the 2018 midterm elections switched their turnout decision 25.9% of the time.

This implies that individuals that are less likely to turn out in an actual election appear to be more likely to be influenced by election forecasts when making a turnout decision compared to individuals that have a historical tendency of voting in elections. Considering that this relationship is only found among “yes to no” switching subjects (which are far more common

than “no to yes” switching subjects), this lends credibility to the aforementioned “survey bias explanation”. There’s a possibility that the wording of the question influenced typical non-voters that were willing to vote in a zero-cost environment to reconsider that decision after being primed to consider the costs of voting and being shown a forecast of the current state of the race.

Of the 186 subjects surveyed, only six switched their turnout decision from “no to yes”. As such, no clear statistical conclusions can be drawn regarding “no to yes” switchers. It is clear however, that “no to yes” switching occurs at a much lower rate than “yes to no” switching, but this is not believed to be a result of any underlying effects of election forecasting. Instead, the number of individuals that intended on voting for their preferred candidate is much higher than the number of individuals that declined to vote for their preferred candidate. Since this base rate is lower than the base rate for those that initially intended on turning out, it’s not surprising that the number of subjects that switched their turnout decision would also be lower, all else equal.

One of the benefits of this study, apart from providing more evidence than there was previously to support the idea that election forecasts do impact voter turnout, is that researchers that conduct similar studies in the future can have a clearer understanding of the required sample size to obtain insightful results. The study covered in this paper does suffer from a low number of “yes to no” and “no to yes” switchers, because the total subpopulation of “switchers” represents roughly 15% of the total population of subjects. As such, it’s recommended that these base rates are considered when conducting future studies, to ensure that enough “yes to no” and “no to yes” switchers can be surveyed. For instance, if this study had obtained a sample size of 1,000, it could be expected that there would be around 150 switchers. With this new knowledge

and insight, clearer conclusions could be drawn regarding the effects of the individual types of forecasts on voters' intentions to turn out.

7. Implications

What's stopping Fox News, CNN, or any media outlet from publishing polls or forecasts on election day that either intentionally or unintentionally deter or encourage voter turnout for a particular candidate? As it stands, nothing is. This may be cause for concern, because the results of this research show the possibility that election forecasts may impact voter turnout. If future research finds that forecasts can impact voter turnout with a substantial effect size² then I worry that this problem will manifest in future elections. It's unlikely that the FEC will attempt to regulate election forecasting and information sharing, but if no restrictions are put in place, how can we trust that our democracy isn't being undermined? Just as fake news has the potential to alter public perception on issues and impact democratic elections, fake forecasts could potentially be used as an instrument by bad actors to siphon votes away from a candidate, or push votes towards a candidate.

One solution may be to continue to improve our forecasting methods so forecasts are more accurate and reliable. This would mean that if people are to use election forecasts as decision aids, at the very least, they're basing their decision off of good and reliable information, so long as they're interpreting the forecast correctly. Media outlets have a moral obligation to verify the credibility of polls and forecasts that they intend to publish prior to publishing them, just as they do in regards to sharing the truth in their reporting of any other story. The truth is

² In this case, a substantial effect size arguably could be anything non-zero, as one vote could be the difference between a particular candidate winning or losing the election.

however, that media outlets misunderstand forecasts frequently, leading them to share information that is unreliable and misleading. This is especially concerning because cable news stations mention forecasts multiple times per day throughout their programming. On average from August 1st, 2016 to Election Day 2016, Fox News mentioned forecasts 3 times a day, CNN mentioned them 2.5 times a day, and MSNBC mentioned them 5 times a day (Messing, 2018).

Additionally, it's grossly irresponsible for media outlets to present election forecasts and polls to audiences without providing important context and properly explaining how to interpret and understand what they mean. As aforementioned, many people struggle to interpret and distinguish between vote-share and win-probability forecasts. Media outlets must understand how forecast reporting can impact election results and do their best to share information with audiences in an unbiased way. It is negligent and irresponsible for media outlets to not verify the credibility of a forecast, but evidence suggests that forecasts shared with audiences tend not to always be credible. I recommend that an independent third-party should be used to grade forecasts, and media outlets should be careful to share forecasts that are above a certain grade to ensure that the public is given the most reliable information.

For polls, one way that this could work is that a consistent set of criteria, such as sample size, poll type (random digit dialing, online survey, etc.), historical pollster credibility, population sampled (registered voters, likely voters, U.S. citizens, etc.), are used to provide a letter grade associated with a poll. This strategy is similar to that of *FiveThirtyEight's* poll grading, which allows them to assign different weights to polls for aggregate forecasts. Media outlets can mention these grades and provide audiences with an idea of how much they should trust forecasts that are shown. This provides two benefits for viewers. First, viewers are able to

easily understand how much they should trust a poll. Second, news outlets will be less likely to share poll results for polls that have a poor grade to begin with, ensuring that viewers are provided with higher-quality information.

Grading election forecast reliability is a more difficult task than grading polls. This is because often, there are a lot more factors and variables that are considered by models that many forecasts are created from. Furthermore, forecasters don't always make their methodology public, and even if it is public, the reliability of a forecast may not always be easy to measure based on an understanding of the methodology alone. Instead, I argue that forecast reliability should be determined based on historical accuracy of election forecasts created by the forecaster, for the particular race type in question. A recommended approach to grade forecaster accuracy equitably, would be basing grades off of metrics such as calibration plots and Brier scores (Boice and Wezerek, 2019). The grading system I propose can be thought of as as a meta-forecast, or a forecast that predicts the expected accuracy of a given forecast.

By creating a standard for media outlets to provide potential voters with high quality information, communicated in a way that is simple to understand, numeracy effects can be minimized, and risks of forecasts biasing turnout in either direction can be also be reduced.

8. Limitations and Future Work

Overall, there's some notable caveats with this study and the findings should be considered with these caveats in mind. The sample size of the study was small, especially when considering the subpopulations of “yes to no” and “no to yes” switchers. These groups, while most important to the purpose of the research, didn't consist of enough individuals to increase the power of the findings. This problem could have been resolved by sampling a larger number of subjects, but due to monetary constraints, this wasn't feasible in this case.

Another caveat is that this study focused primarily on U.S. elections with participants that were U.S. citizens. This isn't necessarily problematic, but it's important to keep this in mind before generalizing the results of this study to other electoral systems. Future work should investigate how turnout is impacted across various electoral systems. An example of existing research that achieves this is a study that looks at the effects of opinion polling on voter turnout in British elections within a parliamentary system (McAllister and Studlar, 1991). These types of studies can help to identify the exact mechanism that causes turnout switching when an election forecast is viewed.

The work covered throughout this paper is focused on presidential elections for the purpose of increasing the number of subjects that would be initially willing to support and turnout for a particular candidate. It is not expected that the results of this study would be substantially different if this research focused on congressional or gubernatorial races, but it's important to consider these limitations before generalizing, especially since the mechanism for turnout switching as a result of viewing an election forecast is currently unclear.

It's also important to consider the fact that the study didn't take place within an actual presidential election. Instead, a hypothetical presidential election scenario was used, with two imaginary candidates, with fictional platforms. The decision to do this was to ensure that the results of legitimate elections were not impacted by conducting this research. It's difficult to confidently know whether or not the behaviors expressed in a hypothetical election would differ from the behaviors expressed in an actual election. It's worth considering the idea of conducting some of the future research in a controlled environment, where study participants are asked to vote in a fictional local election.

The thresholds used for the different forecast types were within the predefined boundaries for "close" and "landslide" races, but the forecasts that were shown to participants were not within the extremes of these boundaries. It's possible that by showing forecasts that give a 5% chance of a candidate winning as opposed to a 20.6% chance, individual forecast-type effects will be heightened. Future research should replicate the study in this paper, but change the actual forecast values in an attempt to increase or decrease the amount of turnout switching that occurs. I hypothesize that these individual forecast-type effects on turnout switching will become more evident if the forecasts shown to subjects are on the extreme ends of what's considered a "close race" and what's considered a "landslide race".

Similarly, it would be interesting to zoom in on particular race types, showing subjects a few different forecasts within that race type, and after each one, asking participants if their initial voting decision would change. In doing this, it would be possible to understand the particular percent chance values that impact individual's turnout decisions, and see if there's a common trend among individuals that decide to switch their turnout decision based on a forecast. This

type of study could also clear up concerns relating to how subjects in this study interpreted the forecasts, and how numeracy plays a role in their turnout choice. If the results of this research were to show that there was no clear pattern to when participants switched from their initial turnout choice, then the numeracy explanation would make much more sense.

In order to rule out the “survey bias” explanation pertaining to the wording of the question that subjects are asked after viewing the forecast, a future study should be conducted where one of two experimental design choices should be made. The first choice would be to avoid asking subjects to think about their typical costs of voting when deciding whether or not to vote. The second choice would be to ask subjects to think of their typical costs of voting when asking them if they’d vote initially and after they’ve been shown the forecast. The benefit of the former option is that there would be no concerns that the questions themselves would have an impact on participant’s turnout decisions. However, in doing this, the costs associated with voting are ignored since the election is hypothetical, which means the results of the study may not be as transferrable to actual elections under this experimental design. The latter option provides the opportunity for subjects to think about the costs associated with voting in an equal way. Subjects are reminded to think about this before both of their turnout decisions. One concern with this approach is that it’s possible for some subjects to read the reminder to think about the costs of voting in only one of the two questions, but this really isn’t too large of a concern. As such, I would recommend that future work takes the latter approach to this problem.

The research in this paper only considers how subjects make decisions when viewing a single forecast on the day of an election. It could be valuable for future research to consider how turnout decisions could be affected when individuals view forecasts over the course of weeks or

months leading up to an election. I hypothesize that forecasts that are not what one would consider “recent” will not be used by individuals as a decision aid for turnout decisions, however, earlier forecasts may be used as a decision aid to continue to view forecasts leading up to an election.

Instead of having subjects view forecasts in the weeks or months leading up to an election, it could also be valuable to see how people interpret and make decisions when shown multiple forecasts simultaneously on the day of an election prior to making their decision whether or not to turn out. This type of study could provide interesting insights regarding how individuals make turnout decisions, as well as shed light on how information or choice overload and numeracy plays a role in decision making processes when forecasts are used as a decision aid. Outside of the lab, we have access to multiple sources and the ability to view more than a single forecast, so understanding how we process multiple forecasts regarding the same event and use that information to make a decision could be insightful as well.

9. Conclusion

This research uses a novel, survey-based, approach in order to understand how voting behavior is affected after potential voters view election forecasts. While the results of this research still leaves many questions unanswered, it is clear that some potential voters' intentions change after being shown an election forecast. Surprisingly, this effect doesn't appear to be associated with forecasted race closeness.

The methodology used in this research provides an effective and promising experimental design for future studies on this topic. "Experimental electionomics", or using surveys to experiment and understand the behavior of voters in hypothetical election environments is an effective tool. It allows researchers to learn about the behavior of voters without fear of potentially impacting the results of an actual election. While experimental electionomics is limited in that it's difficult for subjects to feel like there are any real stakes in hypothetical elections, it serves as an ethical and insightful method to understand voter behavior in elections, so long as the experiment is designed with these caveats in mind.

10. References

- Aldrich, John H. 1993. "Rational Choice and Turnout." *American Journal of Political Science* 37 (February): 246–78.
- Ansolabehere, Stephen D, and Shanto Iyengar. 1994. "Of Horseshoes and Horse Races: Experimental Studies of the Impact of Poll Results on Electoral Behavior." *Political Communication* 11 (4): 413–430.
- Barberis, N. (2013). The psychology of tail events: Progress and challenges. *American Economic Review*, 103(3), 611-16.
- Boice, J., & Wezerek, G. (2019, April 04). How Good Are FiveThirtyEight Forecasts? Retrieved from <https://projects.fivethirtyeight.com/checking-our-work/>
- Cox, Gary W., and Michael C. Munger. "Closeness, expenditures, and turnout in the 1982 US House elections." *American Political Science Review* 83.1 (1989): 217-231.
- Delli Carpini, Michael X. 1984. "Scooping the voters? The consequences of the networks' early call of the 1980 presidential race." *The Journal of Politics* 46 (3): 866–885.
- Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: Harper and Row.
- Endersby, James W., Steven E. Galatas, and Chapman B. Rackaway. "Closeness counts in Canada: Voter participation in the 1993 and 1997 federal elections." *The Journal of Politics* 64.2 (2002): 610-631.
- Enos, Ryan D, and Eitan D Hersh. 2015. "Campaign Perceptions of Electoral Closeness: Uncertainty, Fear and Over-Confidence." *British Journal of Political Science* pp. 1–19.
- Ferejohn, John A., and Morris P. Fiorina. 1974. "The Paradox of Not Voting: A Decision Theoretic Analysis." *American Political Science Review* 68 (June): 525–36.

- Goel, S., Obeng, A., & Rothschild, D. (2015). Non-representative surveys: Fast, cheap, and mostly accurate. In *Working Paper*.
- Gosnell, H.F., 1937. Technical research how accurate were the polls? *Public Opinion Quarterly* 1, 97–105.
- Haspel, M., & Knotts, H. G. (2005). Location, location, location: Precinct placement and the costs of voting. *The Journal of Politics*, 67(2), 560-573.
- Simon, H. A. (1985). Human nature in politics: The dialogue of psychology with political science. *American Political Science Review*, 79(2), 293-304.
- Konitzer, Tobias, et al. “Non-Representative Surveys: Modes, Dynamics, Party, and Likely Voter Space.” *Working Paper*, 2017.
- Kunreuther, Howard, Nathan Novemsky, and Daniel Kahneman. 2001. “Making low probabilities useful.” *Journal of risk and uncertainty* 23 (2): 103–120
- Madonna, G. T., & Young, M. (2017). The First Political Poll. *Accessed December, 20*.
- Marchetti-Bowick, M., & Chambers, N. (2012, April). Learning for microblogs with distant supervision: Political forecasting with twitter. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics* (pp. 603-612). Association for Computational Linguistics.
- Matusaka, John G. "Election closeness and voter turnout: Evidence from California ballot propositions." *Public Choice* 76.4 (1993): 313-334.
- McAllister, I., & Studlar, D. T. (1991). Bandwagon, underdog, or projection? Opinion polls and electoral choice in Britain, 1979-1987. *The Journal of Politics*, 53(3), 720-741.

- Messing, S. (2018, March 22). How election forecasts confuse Americans - and may lead them not to vote at all. Retrieved from <https://www.washingtonpost.com/news/monkey-cage/wp/2018/03/22/how-election-forecasts-confuse-americans-and-may-lead-them-not-to-vote-at-all/>
- Miller, Greg, Ellen Nakashima, and Adam Entous. 2017. "Obama's secret struggle to punish Russia for Putin's election assault." Washington Post .
- Mutz, Diana C. 1997. "Mechanisms of Momentum: Does Thinking Make It So?" *The Journal of Politics* 59 (1): 104–125.
- Mutz, Diana C. 1998. *Impersonal Influence: How Perceptions of Mass Collectives Affect Political Attitudes*. Cambridge, UK: Cambridge University Press.
- Niemi, Richard G. 1976. "Costs of Voting and Nonvoting." *Public Choice* 27 (Fall): 115–19.
- Peters, E., Västfjäll, D., Slovic, P., Mertz, C. K., Mazzocco, K., & Dickert, S. (2006). Numeracy and decision making. *Psychological science*, 17(5), 407-413.
- Pew Research Center for the People and the Press, 14 Oct. 2017. "Political Typology Quiz." www.people-press.org/quiz/political-typology/.
- Riker, William H., and Peter C. Ordeshook. 1968. "A Theory of the Calculus of Voting." *American Political Science Review* 62 (March): 25–41.
- Riker, W. H. (1995). The political psychology of rational choice theory. *Political Psychology*, 16(1), 23-44.
- Sanders, Elizabeth. 1980. "On the Costs, Utilities, and Simple Joys of Voting." *Journal of Politics* 42 (August): 854–63.

- Silver, N. (2018, November 06). 2018 House Forecast. Retrieved from <https://projects.fivethirtyeight.com/2018-midterm-election-forecast/house/>
- Singer-Vine, Jeremy. "Grading The 2016 Election Forecasts." BuzzFeed, www.buzzfeed.com/jsvine/2016-election-forecast-grades
- Squire, P. (1988). Why the 1936 Literary Digest poll failed. *Public Opinion Quarterly*, 52(1), 125-133.
- Taleb, N. N. (2007). Black swans and the domains of statistics. *The American Statistician*, 61(3), 198-200.
- Tuchman, S., & Coffin, T. E. (1971). The influence of election night television broadcasts in a close election. *Public Opinion Quarterly*, 35(3), 315-326.
- Tumasjan, A., Sprenger, T. O., Sandner, P. G., & Welpe, I. M. (2011). Election forecasts with Twitter: How 140 characters reflect the political landscape. *Social science computer review*, 29(4), 402-418.
- Wang, W., Rothschild, D., Goel, S., & Gelman, A. (2015). Forecasting elections with non-representative polls. *International Journal of Forecasting*, 31(3), 980-991.
- Wlezien, C., Jennings, W., & Erikson, R. S. (2017). The "timeline" method of studying electoral dynamics. *Electoral Studies*, 48, 45-56.