# Trust in US Politics

stack DATA STRATEGY


**Abstract**

"In general, do you trust other people?" You wouldn't think someone's answer to such an innocuous question would tell you much about their political views. But it does. In fact, knowing whether someone thinks that other people can be trusted or not, known by social scientists as their level of "social trust", is increasingly central to understanding the choices voters are making across the Western world. In this paper, we examine the variable of social trust as it is related to accurately capturing vote intention in pre-election polls of the 2020 US general election. Modelling social trust based on the American National Election Studies (ANES), we examine how this variable interacts across demographics, turnout, and vote intention. Ultimately, our research shows that high levels of social distrust played a role in the polling error in last year's US presidential election and suggests that pollsters and researchers alike need to consider its influence moving forward.


**What is the problem?**

**2.1 Polling error in the 2020 US election cycle**

In the 2020 US election, pollsters struggled to accurately capture levels of support for the Republican Party - even after making methodological changes following the 2016 election. Polls did not capture Trump's level of support, either at the national level or the state level. The last set of national polls were released the day before the election, these were of likely voters and field work took place in the three days prior to election day. They predicted the following: YouGov had Biden +10%, Ipsos had Biden +7% and 538 (which aggregates polls) gave Biden an 8% lead.[1] Biden's actual winning margin was 4.4%.

Poor national polling was compounded by inaccurate statewide polling. The table below shows polling errors in swing states for some of the leading US pollsters - the polling errors occurred across the industry, affecting our pre-election polling too.

| State | Final Result (R/D) | New York Times | | Fox News | | CNN | | 538** | |
|---|---|---|---|---|---|---|---|---|---|
| | | Prediction | Delta* | Prediction | Delta | Prediction | Delta | Prediction | Delta |
| Pennsylvania | 49/51 | 47/53 | -4 | 47/53 | -4 | 45/55 | -8 | 48/52 | -2 |
| Michigan | 48/52 | 46/54 | -4 | 44/56 | -8 | 43/57 | -10 | 46/54 | -4 |
| Wisconsin | 50/50 | 44/56 | -12 | 48/52 | -4 | 46/54 | -8 | 46/54 | -8 |
| North Carolina | 51/49 | 48/52 | -3 | 48/52 | -6 | 47/53 | -8 | 49/51 | -4 |
| Iowa | 54/46 | 48/52 | -12 | NA | NA | NA | NA | 49/51 | -10 |

1

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| Arizona | 50/50 | 47/53 | -6 | 45/55 | -10 | 48/52 | -4 | 48/52 | -4 |
| Georgia | 50/50 | 50/50 | 0 | 49/51 | -2 | NA | NA | 49/51 | -2 |
| Texas | 53/47 | 52/47 | -2 | 49/51 | -8 | NA | NA | 51/49 | -4 |
| Florida | 52/48 | 48/52 | -8 | 51/49 | -2 | NA | NA | 51/49 | -2 |

***** For 538, their last polling average***

**** Delta here is the difference between the predicted gap and the actual gap***

Figure 2.1: Polling error in the 2020 US Presidential Election

After the 2016 election cycle, many research groups investigated potential reasons for the polling misfire. The American Association for Public Opinion Research (AAPOR) assessed several theories for the fact that the majority of pre-election polls underestimated support for Donald Trump.[2] Ultimately, they found that most evidence suggested an uncaptured late swing to Trump and a pervasive failure to adjust for overrepresentation of college graduates answering polls. In addition, they found little evidence for the *Shy Trump* theory - the theory that voters might lie in polls because of the perceived social undesirability of voting for Donald Trump.

Pew Research also investigated potential reasons for the polling misfire, with the top contenders being a) that pollsters were not appropriately weighting polls to education levels and b) that a section of the Trump vote was being systematically underrepresented in polling samples. Ultimately, Pew's Director of Survey Research concluded that there was little to no evidence of point b - and so the 2016 polling error was attributed to the incorrect application of education weights.[3] These sentiments have been echoed by other academics and political scientists; most commonly, 2016 polling errors have been accredited to a misrepresentative sample or a weighting issue.[4]

In response to this, pollsters altered the way they weighted education to rectify these mistakes for the 2020 election.[5] This weighting adjustment clearly did not have the desired effect, with the polling error in 2020 surpassing the one seen in 2016 - it seems that the systematic underrepresentation of Trump voters which plagued pollsters in 2016 reared its ugly head again in 2020[6]. Given that Trump has recently

---

[2] Kennedy, Courtney, et al. "An evaluation of the 2016 election polls in the United States." Public Opinion Quarterly 82.1 (2018): 1-33.

[3] https://www.pewresearch.org/methods/2021/03/02/what-2020s-election-poll-errors-tell-us-about-the-accuracy-of-issue-polling/

[4] Sturgis, Patrick, et al. "Report of the inquiry into the 2015 British general election opinion polls." (2016); Prosser, Christopher, and Jonathan Mellon. "The twilight of the polls? A review of trends in polling accuracy and the causes of polling misses." *Government and Opposition* 53.4 (2018): 757-790.

[5] https://fivethirtyeight.com/features/what-pollsters-have-changed-since-2016-and-what-still-worries-them-about-2020/

[6] https://fivethirtyeight.com/features/the-death-of-polling-is-greatly-exaggerated/

teased a 2024 presidential run at the 2021 CPAC, pollsters have no choice but to revisit the problem: why can't polls accurately capture levels of support for Trump and what can be done about it?

**2.2 Social trust**

The concept of social trust - the level to which you trust others - is not new in the arenas of social and political science. Political Data Scientist and in-house pollster for the 2012 Obama presidential campaign, David Shor, said in a 2020 interview with the Intelligencer that "working-class people with low levels of social trust were much less likely to answer those phone polls than college-educated professionals" when discussing failures on Hilary Clinton's 2016 presidential campaign.[7]

The argument is this: a person's level of social trust is instrumental in shaping that person's worldview, ideology, and thus political opinion. High trust individuals are demographically and politically distinct from low trust individuals, and in general they lean Democrat. This has two important effects on the sample in any poll. Firstly, demographic groups with low levels of social trust are less likely to respond to polls, meaning that those groups will be underrepresented, skewing a poll towards the Democrats. Secondly, the high-trust individuals *within each group* are also more likely to answer polls than their low-trust counterparts. Thus, even if a poll uses quotas to be demographically and politically representative, it will *still* be biased towards the high-trust, Democrat-leaning individuals in each group. Pollsters made strides to solve the first problem – inter-group selection bias – after the polling error in 2016, but failed to account for the second – intra-group selection bias – and it is this that explains much of the polling error in 2020. To rectify this, pollsters must begin to weight their polls to trust levels, much akin to how they currently weight to age and education levels.

**Modelling Social Trust**

In order to weight a poll to social trust in the United States, one must first find a measure of social trust per state. In other words, in the same way we can say that X% of Texas is white - can we, in some sense, get an accurate measure of the proportion of Texas' population which is 'trusting'?

To do this, we first have to identify both a nationally representative survey and a question related to social trust from which to model from. In this case, we had to turn to external data sources for multiple reasons. First, if we are testing our hypothesis that low trust voters do not participate in polls, then we could not use a self-commissioned voting intention poll that we believed to be unrepresentative in

[7] https://nymag.com/intelligencer/2020/07/david-shor-cancel-culture-2020-election-theory-polls.html

terms of trust. Secondly, in our internal polling, we did not ask a general inter-personal trust question in a nationally representative poll. We therefore had to find a survey whose sample we could be confident had the question we needed and was truly representative.

**3.1 Methodology**

The 'American National Election Studies (ANES) 2016 time series study'[8] presented itself as the clear candidate for a nationally representative survey which asked a key social trust question. This survey consisted of 3,648 respondents, 30% of which completed the survey face-to-face and the remaining over the internet. Respondents

are sampled at random from a selection of addresses provided by the US Postal Service - phone numbers which match these addresses were then obtained from data providers.[9] Before any data is collected, field workers called respondents to verify their address and identity - it is this rigorous process which increases confidence in the representativeness of the sample.

In addition to the high quality nature of the sample, the survey itself includes all required demographic/political questions as well as a question on social trust - it is for all these reasons we used this dataset. The specific question we used on social trust was:

"Generally speaking, how often can you trust other people?"

The possible responses are listed below:

- Always
- Most of the time
- About half the time
- Some of the time
- Never

We transformed this variable into a binary format by assigning all those who answered "always" or "most of the time" as trusting (=1) and all other responses as distrusting (=0) - the resulting data was our dependent variable. To model this, we used a selection of demographic indicators as our independent variables. These were:

- Age (6 levels)
- Gender (2 levels)

[8] 2017. The ANES 2016 Time Series Study [dataset]. Stanford University and the University of Michigan
[9]

- Education (4 levels)
- Ethnicity (4 levels)

The state of residence for each respondent was also used in the model. Once we had our dependent and independent variables, we employed a Bayesian logistic regression model to predict the probability that a person with a certain demography and from a certain state would be trusting. The model was constructed with varying intercepts and multiple varying slopes, with explanation to follow. First we define the probability of the data and the linear model:

$$P_i \sim Binomial(1, p_i) \quad \text{[probability]}$$

$$logit\ (p_i) = state_{[i]} + \alpha_{AGE[i],R[i]} + \beta_{EDU[i],R[i]} + \delta_{GENDER[i],R[i]} + \theta_{ETH[i],R[i]} \quad \text{[linear model]}$$

The linear model for logit(p) contains an intercept term, state(i) which allows general trust levels to vary per state - this is the “varying intercept” feature. The model also considers four separate effects for each demographic in each region. Region is denoted as R[i], in this analysis we used five standard regions of the US to group states: West, South, Southwest, Northeast and Midwest. Each of these four effects are “varying slope” effects, for example, our gender term actually comprises 10 separate coefficients - one term per gender per region (2 x 5). This technique allows for interaction terms to vary by region; trust levels of men in the south and men in the northeast vary in this model.

This is a “partially pooled” model; data is in essence shared to ensure that probability distributions are drawn for each coefficient. The most intuitive example of how this model leverages pooled data is the state term. It is quite common for US polls to contain very few responses from certain states (Wyoming, Kansas for example) and miss some out altogether (Alaska, Hawaii) - by pooling data, these states can still be included in the analysis. If a state is not included in the survey - and thus there is no data to draw a state level term from - the model will draw from the general distribution of all other state level terms and assign a mean state level term. This is also the case with the varying slope terms - this feature of bayesian modelling allows for increasingly complex modelling without the need for massive datasets.

As with any Bayesian model, it is incumbent upon the modeller to specify priors. Below are the priors set for both the state term and the gender term (priors for all varying slope effects are specified in the same way):

$$state_{[i]} \sim Normal(0, \sigma_{state}) \quad \text{[prior for each state]}$$

$$\sigma_{state} \sim Normal(0,1)$$ [prior for standard deviation of states]

$$\begin{bmatrix}\delta_{i,R=1}\\ \delta_{i,R=2}\\ \delta_{i,R=3}\\ \delta_{i,R=4}\\ \delta_{i,R=5}\end{bmatrix} \sim MVNormal\left(\begin{bmatrix}0\\0\\0\\0\\0\end{bmatrix},\ \rho_{GENDER}\right)$$ [prior for gender term]

All priors above are ”non-informative.” We did not build any prior knowledge into our model - in part because we did not have any previous knowledge specific to trust levels, but also because the size of our dataset reduced the need for realistic priors to be set. In simplified terms, our first term denotes that our 51 state level terms (including DC) should be drawn from a normal distribution, whose prior mean is 0 and whose standard deviation is described by σ_state. This standard deviation is itself to be drawn from a normal distribution, with prior mean equal to 0 and prior standard deviation equal to 1. The state term is assigned a 1 dimensional Gaussian prior - each state level term is drawn from a common distribution - as defined by σ_state.

The prior for the gender term is slightly more complex, as not all gender terms are drawn from the same distribution and don't share a common standard deviation. Each varying slope effect is assigned a multivariate Gaussian prior, each prior is five dimensional, because there are five regions. The five sets of gender terms (one for each region) are related through a covariance matrix ρ_GENDER. This means that each gender coefficient (one for male and one for female) is drawn from a distribution that is common to that region, but different to all others. Note that the first term (the matrix of zeros) is simply assigning prior means of zero, something we do because of our model’s inclusion of a state intercept term.

We ran this model in R using rstan.[10] Rstan uses MCMC sampling to draw probability distributions for each coefficient from the posterior (ANES data). The model was specified to run with 4000 iterations and two (Markov) chains. The model ran with zero divergent transitions, Rhat’s (the convergence and efficiency diagnostic term) equal to one which indicated our Markov Chains mixed successfully and sufficient “effective sample size” for all coefficients.

### 3.2 Stratifying to the population

[10] https://cran.r-project.org/web/packages/rstan/rstan.pdf

The model functionally outputs a probability distribution for each of the coefficients specified. As an example, the model outputs 4000 coefficients (equal to the number of model iterations) for the combination of: gender = male and region = 1. We now have coefficient distributions for every possible permutation of the US population (as defined by our demographic and geographic selections). This means that we can apply our model coefficients to the following person:

- Jeff who is 30, has a degree, is white and lives in Florida (which is in the South)

This calculation will return the probability that Jeff would answer our question on social trust with yes. In other words, it gives us the probability that someone of Jeff's exact demography and location is trusting.

In order to determine levels of trust within a state, we have to apply the coefficients from this model onto the US population. To do that we used the 2019 American Community Survey (ACS), which provides detailed tables on the demographic make-up of each state. Combining our model with the ACS data allowed us to predict the percentage of each state which is trusting, based on the demography of that state - this is a process known as stratification. Our final "post stratification" table consisted of 9,792 distinct cells. These cells cover every unique combination of the following: gender (x2), age (x6), ethnicity (x4), education (x4) and state (x51).

Application of the model onto the post stratification table allowed us to rank states from least trusting to most trusting. The top six most and top six least trusting states (defined by their delta to average trust levels) are shown below. One conclusion of this research is already apparent - it is not the case that so-called "red" states are the most distrusting and "blue" states the least. This is the first indication we found that the relationship between trust and political affiliation is more complex than simply, Democrats = trusting, Republicans = distrusting. Therefore, we need to understand how social trust interacts across other variables such as demographics and turnout.

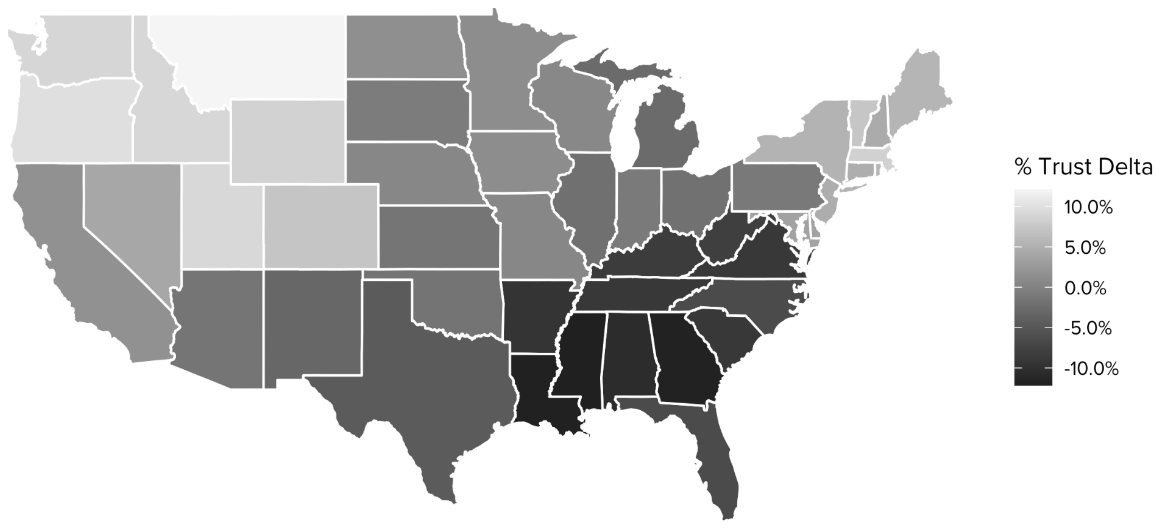


Figure 3.2: Map of trust levels in the United States, modelled using ANES data

**Social Trust and Voting Intention**

**4.1 Is social trust directly predictive of vote intention?**

By plotting our modelled trust against 2016 Republican vote share by state, we see something which again seems counterintuitive - that there is only a slight, almost negligible, negative relationship. Aggregate levels of trust within a state's electorate do not explain levels of Republican support. This demonstrates the increased complexity with the variable of social trust; alone, it cannot be used as an explanatory variable to predict vote intention. This is not particularly shocking, if trust was a direct and singular predictor of vote intention, it would discount the previously explored and finitely proved relationship between vote intention and demographics. Therefore, understanding how trust interacts with demography becomes essential.

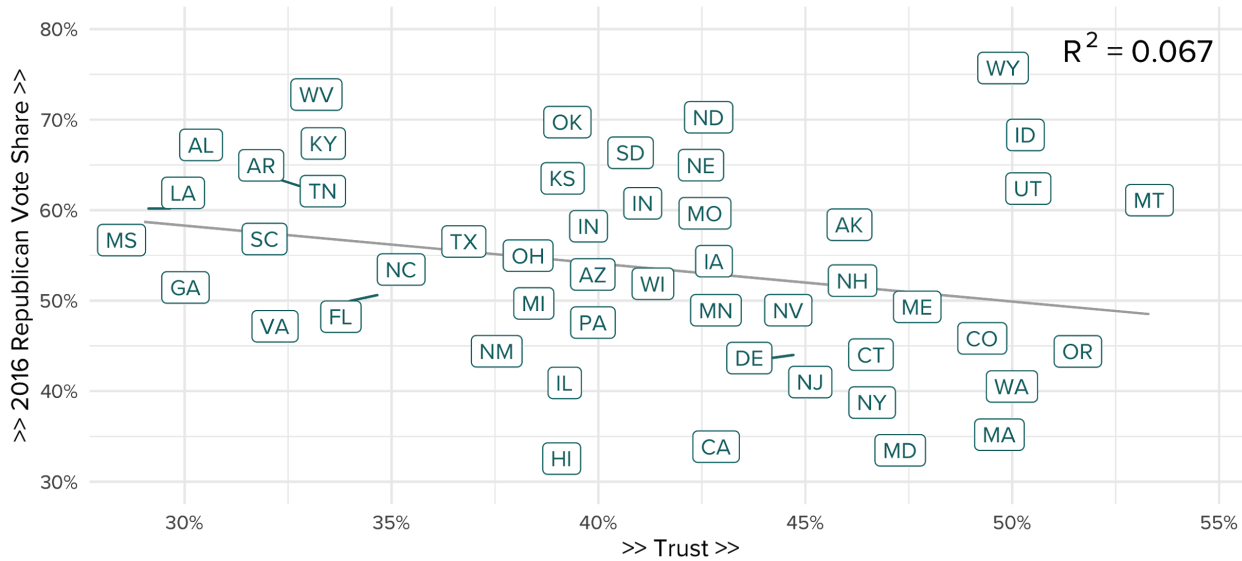


Figure 4.1: 2016 Republican vote shares plotted against modelled levels of social trust

**4.2 Relationship between social trust and demographics**

Trust bifurcates interestingly across various demographic groups, these bifurcations can be seen in the box plots below (Figure 4.2). These visuals were created from modelling which once again used the ANES dataset. Each dot represents the trust level of that demography, red dots represent Republican voters and blue represents Democrats. There are five points per demography and party, one for each region in the US, as outlined in section 3.1.

The interactions between demographics, levels of trust, and vote intention expose multiple areas of interest. Specifically, we are interested in instances where there are

large differences, or deltas, between parties across demographics groups - it is these demographics which we suspect to be behind polling error.

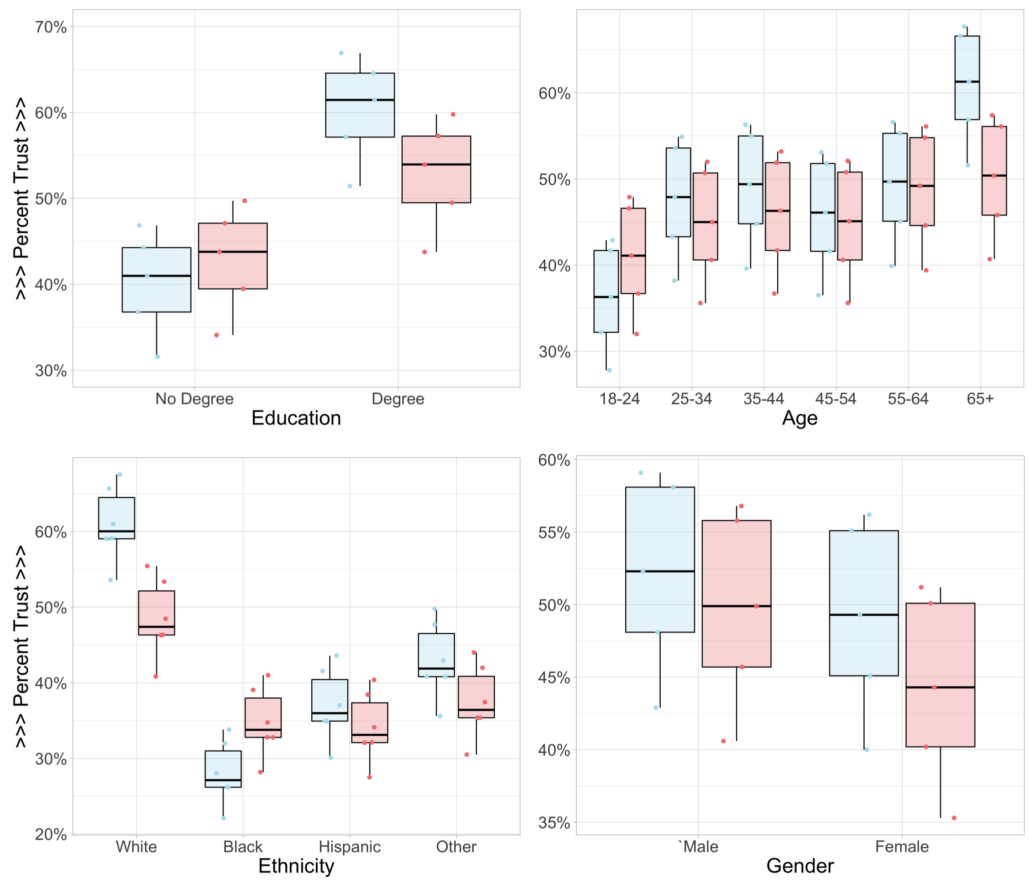


Figure 4.2: BoxPlots of the relationship between trust and different demographics

There are multiple groups of interest which require a level of commentary within the framework of our thesis:

1. *Degree educated Democrats vs Republicans*
   Democrats with a degree are on average 8% more trusting than their Republican counterparts. To give the starkest example, degree educated Democrats in the south west (Arizona, Texas, Oklahoma and New Mexico) are a staggering 23% more trusting than degree educated Republicans in the south.

2. *Old Democrats vs Old Republicans*
   When considering the age plot, the most striking difference is for those who are 65 or older. Republicans in this age category are on average 10% less trusting than Democrats. You can already start to see that within the context of this hypothesis, a quota defined on age - but not accounting for trust - will tend to capture many more Democrats than Republicans.

3. *White Democrats vs White Republicans*
   One of the most striking differences in all of the above plots is the difference in trust levels between white Democrats and white Republicans. White Republicans are on average 13% less trusting than their Democrat counterparts. For example, Iowa and Wisconsin, from which came some of the most disastrous polling of the cycle, have electorates which are 91% and 86% white respectively, highlighting the issue.

4. *Black Democrats vs Black Republicans*
   A curious difference and one that deserves commentary is that between black Republicans and black Democrats. Black Republicans are 6% more trusting than black Democrats. Combine this with the fact that young Democrats are less trusting than young Republicans and you start to build a hypothesis which explains the small but significant swings from Democrat to Republican in this demographic in 2020. Again, polls would struggle to uncover this, as young, black (previous) democrats would not be captured in polling.

The potential impact of polling not accurately capturing these groups in part depends on whether or not they actually vote. If a group is not captured by polls but doesn't vote anway, then they are not responsible for polling error. Combining the above with turnout rates is the next chapter of our analysis.

**4.3 Turnout Model - Methodology**

At this point, it is worth revisiting our thesis within the context of findings within previous sections.

1. Levels of social trust cannot, at the aggregate level, determine how a state will vote (Figure 4.1)
2. A person's level of social trust is strongly related to their demography and political affiliation (Figure 4.2)

But how does this impact polling? To begin to analyse this question, we must switch our focus from levels of social trust within the population, to levels of social trust within the electorate. A low trusting individual who is missing in polls but is also a part of the 30% of people who do not vote will not contribute to polling error - we are solely concerned with those who vote, but whose levels of social trust mean that they do not trust and thus do not take polls.

In order to analyse the electorate, we had to transform our post stratification table with a turnout model. Modelling turnout is a notoriously difficult thing to do, mostly because self reported turnout - that is, asking someone if they voted and taking their word for it - tends to be incorrect, people often say they voted when they didn't. The only way to guarantee that a person did indeed vote is to match their name and address to voting records. Such a dataset, which asks for turnout and validates responses, exists within the ANES' 2016 data collection. Their "Vote Validation" dataset contains unique IDs which can be matched with the previously used 2016 Time Series dataset. This combined dataset contains all demographic information needed for modelling, as well as political information (2016 vote) and validated turnout.

Before modelling turnout, we performed some manipulations on our combined dataset: we mapped all demographics (consistent with mapping for the social trust model), our previous vote and transformed validated vote into a binary variable, where one meant that the respondent definitely voted in the 2016 Presidential Election. In a similar process to the one outlined in section 3.1, we employed a Bayesian logistic regression to model turnout, with the exception that this time, our model included a hierarchical term. As in section 3.1, we first define the probability of the data and the linear model:

$$P_i \sim Binomial(1, p_i) \qquad \text{[probability]}$$

$$logit\ (p_i) = state_{[i]} + \alpha_{AGE[i],R[i]} + \beta_{EDU[i],R[i]} + \delta_{GENDER[i],R[i]} + \theta_{ETH[i],R[i]} \qquad \text{[linear model]}$$

However, our state term is now hierarchical - we define it as such so that it is modelled independently of our linear model. This term allows us to include additional informative data in our model. Our hierarchical state term is defined as:

$$state_{[i]} \sim Normal(\mu_{state}, \sigma_{state}) \qquad \text{[hierarchical state term]}$$

$$\mu_{state[i]} \sim \left(\gamma_{state[i]} * \sigma_{state}\right) + (\tau_{turnout} * turnout_{[i]}) \qquad \text{[state term definition]}$$

$$\sigma_{state} \sim Normal(0,1) \qquad \text{[prior for standard deviation of states]}$$

$$\gamma_{state[i]} \sim Normal(0,1) \qquad \text{[prior for each state]}$$

$$\tau_{turnout} \sim Normal(0,1) \qquad \text{[prior for turnout]}$$

The external data feeding into this term - shown above as "turnout" - refers to the percentage turnout for each state from the 2016 General Election. Functionally, this term acts to increase turnout predictions in states where historic turnout is high and depress turnout where historic turnout is low. All priors for this model were set identically to those outlined in section 3.1. Coefficients from this model were then applied to each cell within our post stratification table, returning the rates at which each individual cell voted in 2016.

In order to make our post stratification table truly representative of the most recently known electorate - the 2016 electorate - we also modelled 2016 vote into it. We used an identical model and the same input data as previously mentioned, but used previous vote as our dependent variable. This model also differed from the turnout model as it's hierarchical term included 2016 Republican/Democrat vote shares instead of turnout.

By applying this party model to our post stratification table, we were able to break-down the US population into 19,584 distinct cells - 9,792 (half) of which were cells which voted Republican in 2016 and the other half Democrats. It is worth mentioning, without exploring in detail, that this post stratification table was then "raked" - this is the process by which we ensure that at the aggregate level, our post stratification table matched known marginal distributions. In this case, after raking, our aggregated table included the precise number of Republicans and Democrats in each state from 2016 and overall the correct number of voters (from a two party perspective) in each state.

### 4.4 Turnout of the US Electorate and Trust

In order to gain insight from our post stratification table, we first had to condense it down from its 19,584 cell form. To do this, we grouped demographics into easily identifiable blocks and plotted their levels of trust against their turnout rates. Combining the Democrat and Republican electorate, Figure 4.2.1 plots these blocks. The size of each bubble represents the overall proportion of that block within the electorate and for ease of viewing, the top 10 largest blocks are coloured green.

This was the first indication that a clear majority of the US electorate are more distrusting than trusting - there are more bubbles below the x-axis than above. It also builds upon analysis outlined in section 4.2, there are large demographic blocks within the electorate which have low trust but high turnout, specifically: old, white, non-degree educated people.

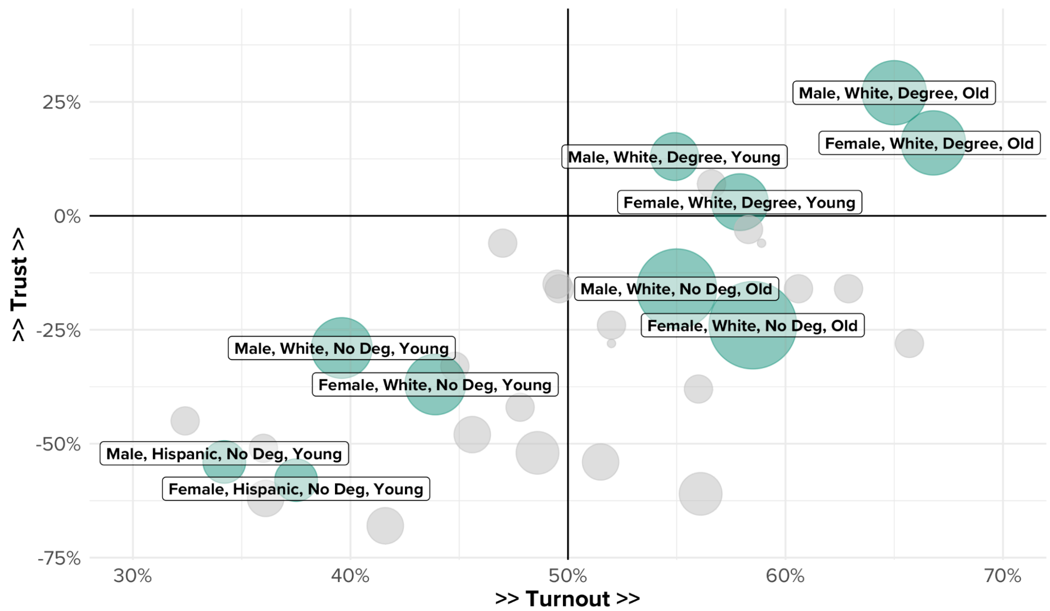


Figure 4.4.1: Trust vs Turnout for the US electorate

Due to the fact that previous vote was included in our post stratification table, we could now plot the same chart as above, but for the Democrat and Republican electorate specifically. The charts below culminate all previous research and in our view, provide compelling evidence of the social trust thesis. One of the most striking findings is that the entire Republican electorate have higher levels of distrust than trust; this is directly contrasted with the Democrat electorate with multiple large blocks being net trusting.

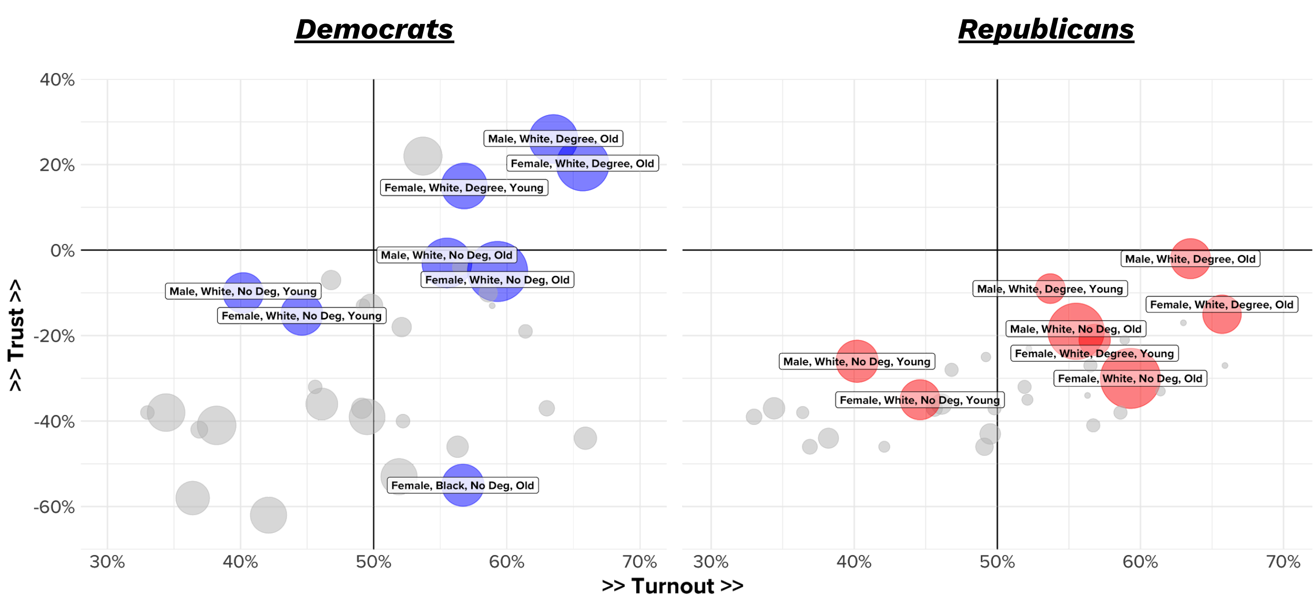

Figure 4.4.2: Trust vs Turnout for the 2016 Democrat electorate (left) and the 2016 Republican electorate (right)

To summarise the above, low trust amongst Republicans occurs in high turnout groups, whereas low trust amongst Democrats occurs in low turnout groups. If you miss a low trust Democrat, there's a good chance they won't vote anyway - but by failing to capture a low trust Republican you're failing to accurately capture the full electorate. Functionally, this finding presents two issues with polling, one specific to the methodology behind "quotas" and the other specific to how polls weight data after collection:

1. Filling Quotas
   When a poll is commissioned, a set of quotas are defined to ensure the sample of people who take the survey comprise a nationally representative group. Most commonly, these quotas specify demographics like age, gender and education - this ensures, for example, that your sample has a nationally representative number of degree vs non-degree educated people. A mechanical issue that pollsters have is filling each quota, this means for example, ensuring that you have approximately 50% male respondents and 50% female. When we start interlocking variables, as many pollsters do, this becomes difficult - ensuring that 7.5% of your sample are degree educated women who are aged 65 and over - is difficult. This is why pollsters "weight" their polls, to correct for the impossible task of finding truly representative samples.

   Pollsters often lift their quotas at the end of field work in order to get the required total sample - a mechanical fact that is important to note within the context of trust. The industry already uses panels of people who are unrepresentative in terms of trust, and then worsens this by filling quotas with people they have remaining in their panel. A poll being unrepresentative in terms of trust is compounded by the fact that the most trusting groups of all are white, college educated Democrats - when a poll under-indexes on low trusting Republicans, it has to over-index on something else, and there are plenty of trusting Democrats in panels to fill quotas with.

2. Weighting to high trust Republicans
   Most voting intention polls will quota on 'previous vote' in order to ensure that in general, the poll captures more Democrats in California and more Republicans in Alabama. After data collection, pollsters apply weights to their data, one of which is a weight to previous vote, correcting for that fact that you cannot ensure your poll captured the exact percentage of Republican/Democrat voters from the last election in each state. It is easier to capture Democrats than Republicans, owing to the fact that the Democratic electorate is more trusting overall - this is an issue, but one that is addressed

by weighting. However, the added trust dimension complicates the issue. The already low number of Republicans which a poll does manage to capture have higher trust levels than the general Republican electorate - they are not representative of the demographic and political groups they belong to. This relatively small number of high trust Republicans are then weighted up significantly, further skewing the vote intention of the Republiucan electorate and hence the outcome of the poll.

### 4.5 Application to other Western Democracies : EU Referendum

It would be one thing if this issue were limited to the United States, but our research suggests that it may be behind the major polling misses in the UK over the past five years and have consequences for politics across the Western world.

In particular, we've found that social trust was a central divide between Leave and Remain voters in the build up to the 2016 EU referendum - another election where the polls missed the mark. Using data from the British Election Study (BES), we can measure the level of social trust among UK voters in the year of the referendum. The social trust question used from the BES was almost identical to the one used from the ANES and outlined in section 3.1. Our UK analysis shows that, as in the United States, the groups central to Leave's victory, white voters without a degree, combine low trust with high levels of turnout.

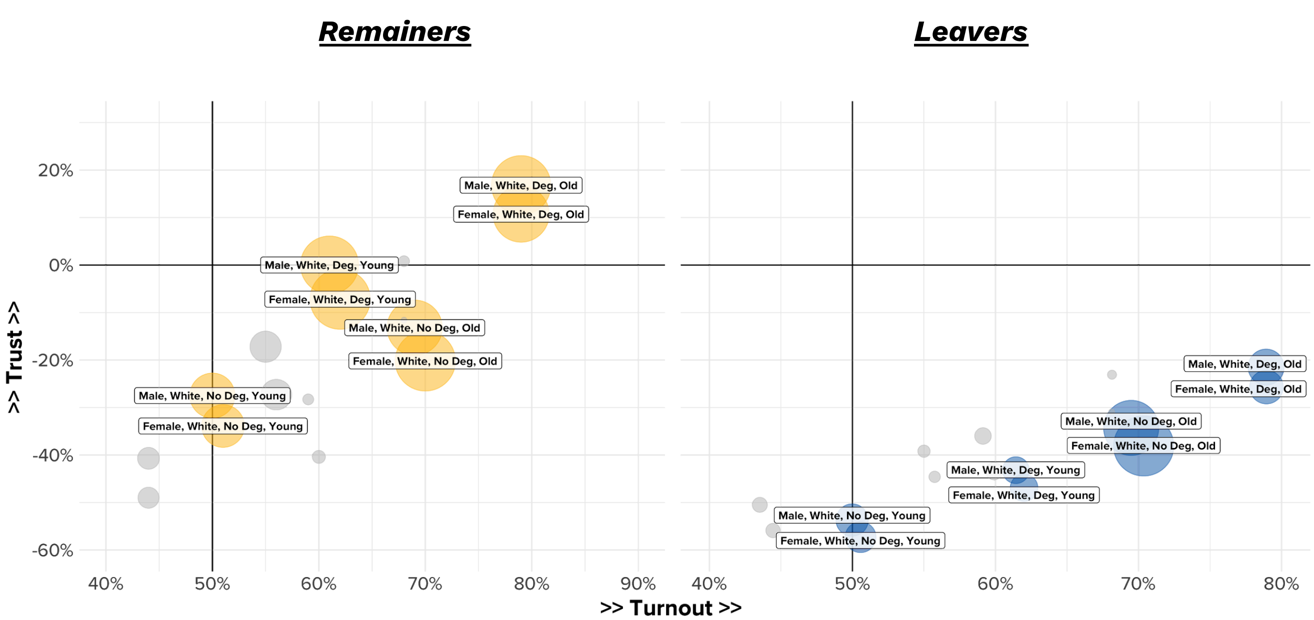


Figure 4.5: Trust vs Turnout for the 2016 Democrat electorate (left) and the 2016 Republican electorate (right)

The other interesting comparison is between the Remain and Democrat electorates. Both have large sections which have high levels of trust. Just as in the US, these voters were numerous in UK panels and hence over sampled to fill the under indexing of sections of the Leave electorate. The same groups drove polling error in both the US and the UK - and these groups are only distinguishable by their levels of trust.

**Employing Trust to Predict Vote Intention**

All of the aforementioned complexity and relationships with the variable of trust has demonstrated its potential influence in predicting vote share in pre-election polling. As trust clearly interacts with demographics (Figure 4.2) and demographics turnout at different rates (Figure 4.4.1, Figure 4.4.2), it is important to consider when assessing voting intention. Just as it is commonplace for pollsters to weight polls to education and age, we believe that including trust as a weighting factor will result in better vote intention results.

**5.1 Weighting to Trust : Methodology**

As our pre-election poll did not include a specific question on interpersonal trust, we had to use an imperfect, but tightly correlated predictor of trust. Our internal poll asked participants about their levels of trust towards the electoral process, which we then weighted to the ANES modelled trust to arrive at our final result.

Our internal poll was conducted in late October 2020 in the lead up to the presidential election. This poll had a sample size of 13,000 respondents across 42 states and asked a range of political and social questions including previous voting behavior, whether they were planning to vote and if so, for who, as well as a range of demographic questions. When weighting to demographics and social trust by state, we were faced with challenges due to smaller sample sizes per state. This was rectified by grouping similar states together in a few cases.[11] For example, in order to reach a sample large enough to weight by all of the mentioned attributes, we had to pool or group some of the midwestern states together to account for a low sample size in Wisconsin.

We weighted our data using the ‘anesrake’ package in R.[12] This package allows the user to rake data using specified target weights which were taken from our modelled post stratification table. The weights used were: age, education, gender, ethnicity, previous vote and trust.

---

[11] Note: we only pooled states in two instances; first, due to the extremely low sample size in WI, we grouped it with Midwestern states. Secondly, there were six states in which the sample were so low that they were unable to be weighted by all demographics (because they didn't include participants from each demographic group) therefore we also grouped IA, WY, ND, SD, MT, AR

[12] https://cran.r-project.org/web/packages/anesrake/anesrake.pdf

## 5.2 Results by State

Employing the methodology outlined in this research reduced the average error on the Trump vote in swing states from 3.0% (in our standard methodology) to 0.9%. On a two party basis, our standard methodology had an average error of 6.0% whereas our trust methodology produced an error of 1.8% - a significant improvement and reduction in error of 4.2%. Our predictions of Trump's vote share based on this trust methodology are shown below, alongside our standard methodological results, and the last poll from major pollsters before the election.

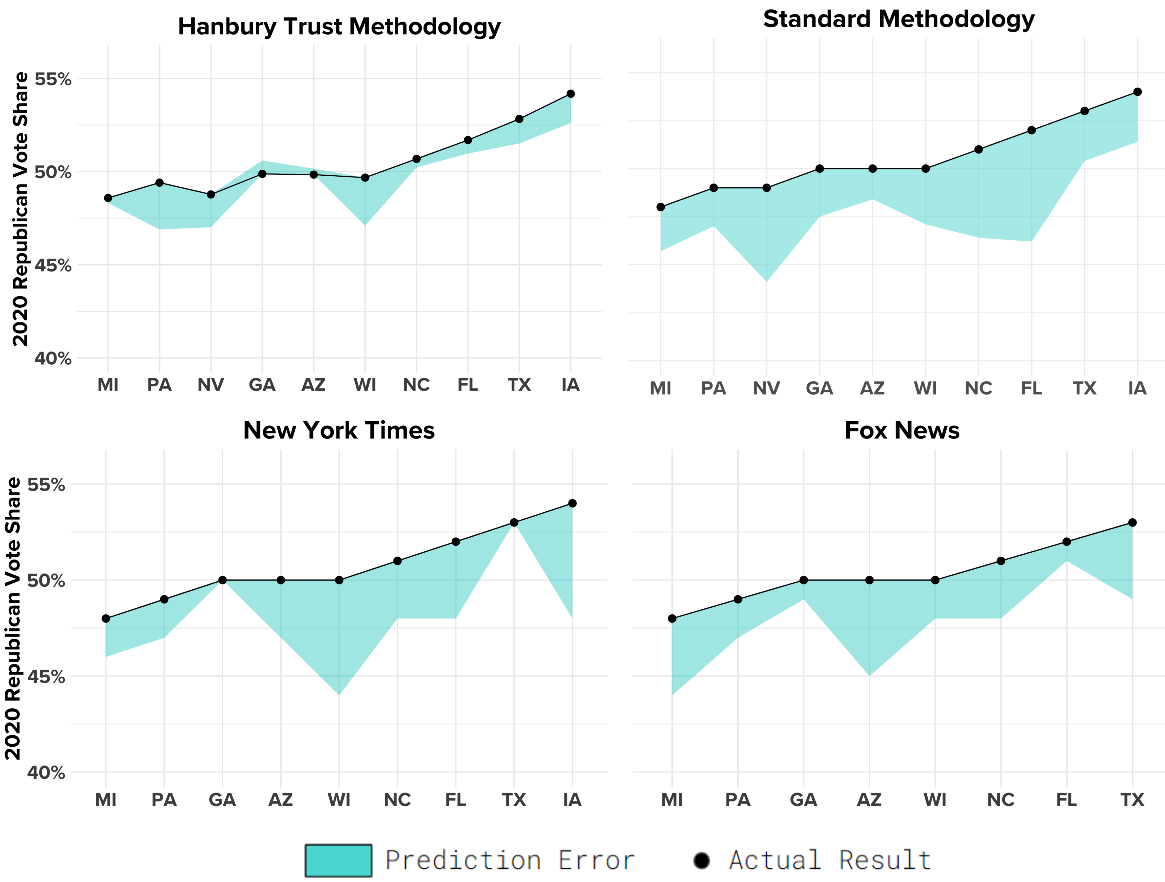


Figure 5.1: Republican vote share estimates in key swing states at the 2020 presidential election

A more in-depth view of our results is presented in the table below. Substantial reduction in errors were seen across all states except Pennsylvania and Wisconsin, where we still predict Joe Biden to win with a much larger margin than he actually did. In both states, we believe that larger and more representative samples would improve the outputs.

One state which requires special attention is Georgia, a state for which polling across the board was excellent. Georgia's black population is a little over 30%, it is the state with the second highest percentage of black voters in the US (second only to Mississippi) and was the centre of a black voter drive in 2020 directed by Stacey Abrams. It is one of a handful of states where distrusting Democrat groups become electorally significant. In general you could classify black Democrat voting blocks as low trust and low turnout, older black women without a degree have the highest turnout of all groups of black voters and can be classified as low trust, medium turnout (Figure 4.2.2). These voters are a significant block in Georgia, in 2016 for example, the black vote accounted for 57% of the Democrat vote share. Within the context of this research, the theory of what happened in Georgia is as follows: polls failed to capture distrusting Republicans – an electorally significant group – as they did in every state. Polls also failed to capture distrusting Democrats who were being encouraged to vote en masse – thus becoming more electorally significant. The errors introduced by missing both of these groups counteracted each other, two wrongs made a right, and pollsters correctly called Georgia.

| State | Final Result (R/D) | Standard Methodology | | Trust Methodology | |
|---|---|---|---|---|---|
| | | Prediction | Delta | Prediction | Delta* |
| Pennsylvania | 49/51 | 47.0/53.0 | -2.4 | 46.9/53.1 | -2.5 |
| Michigan | 48/52 | 45.7/54.3 | -2.9 | 48.3/51.7 | -0.3 |
| Wisconsin | 50/50 | 47.1/52.9 | -2.6 | 47.1/52.9 | -2.6 |
| North Carolina | 51/49 | 46.4/53.6 | -4.3 | 50.2/49.8 | -0.5 |
| Iowa | 54/46 | 51.4/48.6 | -2.8 | 52.6/47.4 | *-1.6* |
| Arizona | 50/50 | 48.4/51.6 | -1.4 | 50.2/49.8 | +0.3 |
| Georgia | 50/50 | 47.5/52.5 | -2.4 | 50.6/49.4 | +0.7 |
| Texas | 53/47 | 50.4/49.6 | -2.4 | 51.5/48.5 | -1.3 |
| Florida | 52/48 | 46.2/53.8 | -5.5 | 51.0/49.0 | -0.7 |
| Average | | | **-3.0** | | **-0.9** |

**** Delta here is the difference between the predicted and actual Trump vote***

Figure 5.2: Predictions for swing states using trust methodology

## 5.3 Implications

### Limitations

Although these results prove to be extremely exciting for the future of polling and predicting voter behaviour, we do recognise the fact that this analysis is not without limitations - in an ideal situation we would have weighted interpersonal trust to our modelled trust values.

In addition, the sample we had access to was not particularly representative of state populations. As we were using repurposed polling, the data was not ideally quota'd to the electorate. Beyond the representation within our sample, we also suffered from a relatively small and disproportionate sample size. Our total sample was 13,000 respondents, but in some key states the sample size was very low - in Wisconsin for example, we only had 240 respondents. In order to prove the impact of trust we need to test this methodology on a more representative and larger sample.

### Conclusion and Future Research

This work emphasises the complex role that interpersonal trust plays in predicting vote intention. We would argue that traditional polling methodologies (both online and telephone) are systematically missing a politically significant part of the population. This not only has consequences for election forecasting but also for market research moving forward.

Following the 2016 election, researchers and political scientists staunchly denied the presence of a *Shy Trump* voter, instead pointing towards sample and weighting concerns. Adjustments were made to address these shortcomings, but the predictions remained skewed to the Democratic Party in 2020 and sampling concerns persisted. We have found that Trump voters may ultimately be *shy* - not in their lack of admission of vote intention - but in their lack of willingness to respond to a poll. Throughout two election cycles, polls vastly underestimated Trump's popularity due to the fact that low-trusting individuals were not appropriately captured in nationally representative polls. By weighting to levels of trust in addition to demographics and previous vote, we were able to predict the actual result more accurately.

This also has broader implications for the future of market research and challenges many established methodologies, like sampling and weighting, which are both commonplace in the industry. If polling in the US has not been capturing a significant proportion of the population for the entirety of Trump's term, then associated findings are most likely not representative of the population at large - instead, such findings represent the views of the "panel-population", a subset of the population defined by its increased trust levels. Such problems regarding polling representation are not confined to the field of election prediction, but should raise concern across other industries. Governments, think tanks, NGOs and corporates all rely on polling - largely from a similar pool of providers and thus a similar pool of panels - to shape strategy, decision making and policy. Consider the scale of the error in predicting the 2020 election and then imagine the impact transposing such an error would have on a range of industries - a Government trying to assess the popularity of a policy, a corporation trying to understand and shape their reputation, a think tank trying to prove a thesis. As polling and the business of understanding public opinion grows ever larger, so does the risk that decisions and directions are based upon findings derived from a sub-section of society.

This methodology needs to be replicated across a range of different polls to truly understand the complex relationship between trust and voting in the Western world. Within Stack Data Strategy, we will be including an interpersonal trust question in our standard demographic collection of questions on all future polls. This will allow us to test our above hypothesis and further refine the methodology across a wide range of polls.