

# Placer County’s Turnout is Not Anomalous

Justin Grimmer<sup>a</sup>

<sup>a</sup>Democracy and Polarization Lab, Political Science, and Hoover Institution. Stanford University.  
616 Jane Stanford Way, Stanford CA 94305. jgrimmer@stanford.edu.

June 12, 2023

**Executive Summary** Douglas Frank’s analysis of Placer County’s turnout data, and his analysis of U.S. elections more generally, is deeply flawed. Frank never demonstrates that his methodology is able to accurately identify systematic voter fraud. Rather, his conclusions are based on mere assertions about his personal view of “uncanny” or “unnatural” relationships. None of these speculations are based on any literature, analysis of historical elections, or mathematical derivations. I will show in this memo that the supposedly unnatural relationships he identifies are the result of profound and basic statistical errors. Frank claims that there is a single turnout rate for partisans across Placer County precincts, but he uses the wrong evidence to test this claim. The correct test shows that Frank’s claim is false, even when using Frank’s own faulty data. Overall, Douglas Frank’s analysis only shows the unsurprising fact that groups with more people tend to have more people turnout to vote. In other words, Frank’s unnatural correlations arise from correlating a variable with itself. Frank attempts to reply to this critique with a nonsensical “bathtub” analogy. This analogy fails to address any of the mathematical or statistical arguments I use to demonstrate that his methodology is flawed and assumes a conclusion he never demonstrates. Frank’s methodology continues to be useless for assessing voter fraud.

Douglas Frank claims to provide evidence of voter fraud based on analyses of turnout data. My research group has documented extensive errors in Frank’s claims. In response to this work, Frank recently produced the first written report of his analyses of U.S. elections (Frank 2023), specifically responding to a memo I wrote regarding fundamental flaws in Douglas Frank’s analysis of elections in Placer County, California.

In this brief memo, I demonstrate that the conclusions in Frank’s report are deeply flawed and that he fails to meaningfully address any of our extensive critiques of his work. More specifically, I will show:

- 1) Frank’s analysis is based on the false premise that predictability implies manipulation. Frank fails to demonstrate that his methodology can accurately identify voter fraud. Frank’s assertions about “unnatural” or “uncanny” relationships is based on his own personal impressions rather than any relevant literature, analysis of historical election results, or mathematical derivations.
- 2) Frank’s own data demonstrates that his claim about a single partisan turnout rate in Placer County is false. Frank used the wrong evidence to assess his claim in his original presentation and repeats that error in his report. When the correct test is used, there is clear evidence of variation in precinct-level turnout rates in Placer County. This conclusion holds regardless of how the precincts are defined, including if I use Frank’s incorrect Placer County precincts.
- 3) Frank’s claim that he can perfectly predict turnout rates in counties is false. Frank uses the wrong evidence to evaluate this claim and incorrectly evaluates predictions. Because of these errors, Frank fails to recognize that his analyses merely demonstrate the unsurprising fact that groups with more people tend to have more individuals who turnout to vote.
- 4) Frank’s use of a “bathtub” analogy to defend his methodology fails to even remotely address any of these critiques. It assumes a conclusion that Frank never demonstrates.

Frank’s analyses are useless for identifying election fraud and provide no evidence of systematic vote manipulation.

## **Frank Never Demonstrates Predictability Implies Manipulation**

Douglas Frank asserts that “By sampling a few counties (or precincts), I have found that I can predict all the rest, providing strong statistical evidence that our elections are being manipulated by computer algorithms” (Frank 2023, 11). Elsewhere in his memo, Frank

asserts that a correlation is “unnaturally high, providing statistical evidence that ballots are being stuffed in Placer County elections according to a mathematical algorithm” (Frank 2023, 8). And Frank claims that “The fact that Placer County registrants ‘vote’ so consistently is strong evidence for central monitoring and subsequent manipulation of their elections (primarily ballot harvesting and stuffing enabled by networked election systems)” (Frank 2023, 10).

Frank provides no evidence for these claims. These are merely impressionistic assertions from Frank, not principled conclusions based on scientific analysis. Frank never shows that his methodology can accurately identify when fraud does or does not occur in U.S. elections. Frank never shows that he only finds correlations of this magnitude in elections where there was known to be manipulation according to a “mathematical algorithm” or “computer algorithm”. Frank does not even provide evidence that the supposedly anomalous correlations occur more frequently in locations where there was known fraud. Nor does Frank show that the correlations and predictability he observes in the 2020 U.S. election are historically anomalous. And Frank fails to cite a single academic study to justify the conclusion that the correlations he observes are consistent with manipulation in U.S. elections.

Frank makes a series of claims about these models that are not based on any statistical or social science literature. Instead, they are based on his own impressions of what constitutes a surprising result. These impressions are false. For example, after describing correlations, Frank asserts that “People familiar with statistics understand the significance of these unnatural  $R^2$  values” (Frank 2023, 2). And he claims that his methodology “can tell you with uncanny accuracy how many of them are recorded as having voted” (Frank 2023, 2). Despite Frank’s claim, there is no statistics or social science literature to justify claims that the correlations he describes are “unnatural”. This is because the extent of correlation varies considerably across different contexts and depends on what is being correlated. As I explain below, the correlations he observes are quite natural. And it is well known that  $R^2$  and predictive accuracy are distinct concepts (Grimmer et al. 2022).

## Frank's Methodology to Assess Placer County Turnout Uses the Wrong Evidence

Frank first presented his analysis of Placer County at the “Moment of Truth Summit”, where Frank claimed to have clear evidence of manipulated results: every precinct in Placer County, according to Frank, had the same rate of Republican turnout. Frank asserted that “88% of the registered Republicans in EVERY Placer County precinct voted...Exactly 88% of all Republicans voted in that Precinct”. After asserting that this is unbelievable, Frank presented a slide where he opines that “we might suspect that someone was stuffing ballots up to a target limit of 88%.” Frank makes a similar claim in his report, asserting that “In order to attribute this to a natural phenomenon, one would have to believe that partisans in every Placer precinct voted at a nearly identical rate; that voter turnout is the same regardless of whether a precinct is rural or urban, minority or diverse, wealthy or poor, or suburban or farming communities. Every Republican in every precinct voted with the same propensity (88.3%) – and so did every Democrat (87.8%). Even voters registered as ‘other’ are unnaturally consistent, they just turned out at a lower rate (76.6%)” (Frank 2023, 9).

Frank uses the wrong test to assess this claim. Frank's evidence presented at the “Moment of Truth” summit examined the relationship between the *number* of Republicans who voted and the *number* of Republicans registered across precincts in Placer County. He also added a line of best fit, called a regression line, for the relationship between the number of votes and the number who are registered to vote. In Frank (2023) he continues to assess claims about constant turnout rates by plotting the number of voters in a precinct against the number of registered voters, along with a line of best fit, as evidence for his claim of a “nearly identical” turnout rate.

This is the wrong test for his claim. As I explain in Grimmer (2022), the plot and regression line merely demonstrate the unsurprising fact that precincts with more registered voters tend to have more individuals who turnout to vote. Differences in the turnout rate are difficult to observe from this plot because differences in turnout rates across precincts

will be obscured by differences in precinct size.

If we want to test a claim about turnout *rates*—for example, to test Frank’s stated claim that exactly 88% of registered Republicans voted in each precinct—we should directly examine turnout *rates*. The appropriate evidence to test claims about “nearly identical” Republican turnout rates in Placer County precincts is to examine the variability in Republican turnout rates across Placer County precincts. This is the evidence I present in Figure 1. There, I examine Republican turnout rates under three definitions of Placer County precincts: the left-hand plot shows the precinct-level Republican turnout rate for the official Placer County data (meaning this plot replicates the right-hand plot in Figure 1 from Grimmer (2022)), the center plot shows the precinct-level Republican turnout rate using Frank’s data which incorrectly divides Placer County into 35 precincts, and the right-hand plot uses the division of Placer County into 538 precincts that Frank discusses in his memo.<sup>1</sup>

Figure 1 shows that Douglas Frank’s characterization of the precinct-level turnout rates in Placer County is false: there is variation in the precinct-level Republican turnout rates. Each figure clearly shows that there is variation in the turnout rates. One way to assess variation is with a standard deviation, which summarizes the average deviation from the mean. Using the correct Placer County precincts, I find a standard deviation in the precinct-level Republican turnout of 13.0%. Using Frank’s own numbers based on the incorrect Placer County precincts I find a standard deviation in precinct-level Republican turnout of 3.9%, and using the more granular precinct divisions I find a standard deviation of 17.3%. In Figures 2 and 3 I show there is variation in the precinct-level Democratic turnout rate (Figure 2) and the precinct-level turnout rate among individuals who do not register as either a Democrat or Republican (Figure 3). Using the correct test, examining turnout rates reveals that Frank’s assertions about a single turnout rate are false.

---

<sup>1</sup>In Douglas Frank’s report he provides the data for all 35 precincts in Placer County. I use those numbers in order to assess his claims. Frank’s precincts appear to arise from aggregating by “Precinct Name” in the voter file. But “Precinct Name” merely provides information about the general location of voters within Placer County. The precinct division I used corresponds to Placer County’s official precincts. The 538 precinct division corresponds to a more granular division, which I obtained from the California redistricting database.

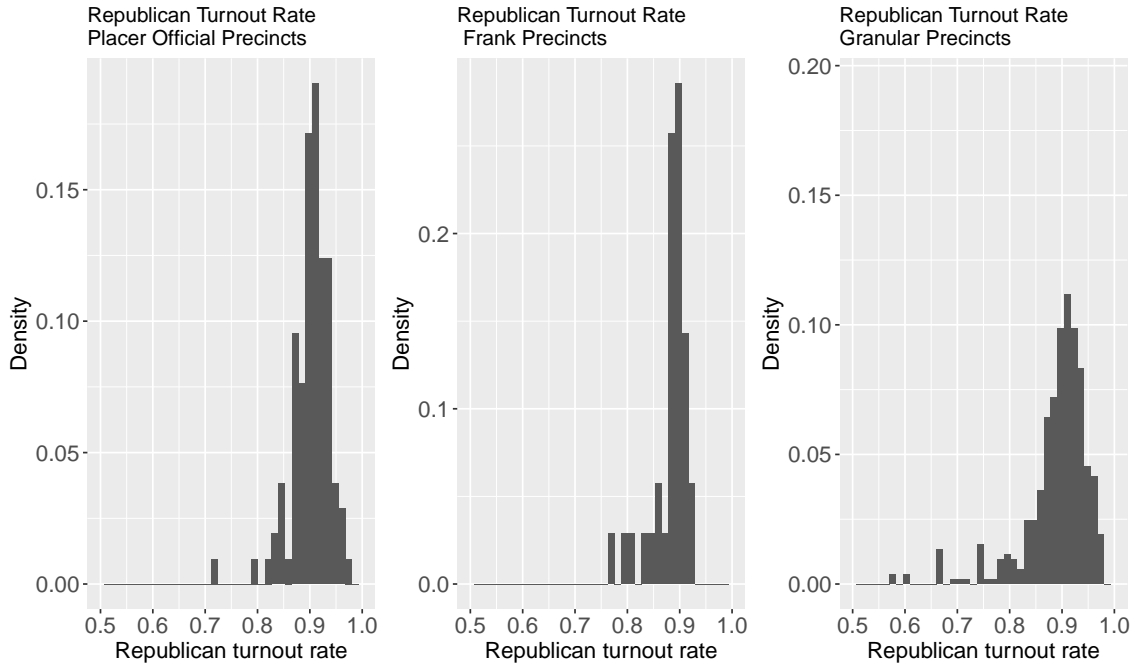


Figure 1: There is considerable variation in the Republican turnout rate across Placer County precincts, regardless of how those precincts are defined.




Figure 2: There is considerable variation in the Democratic turnout rate across Placer County precincts, regardless of how those precincts are defined.

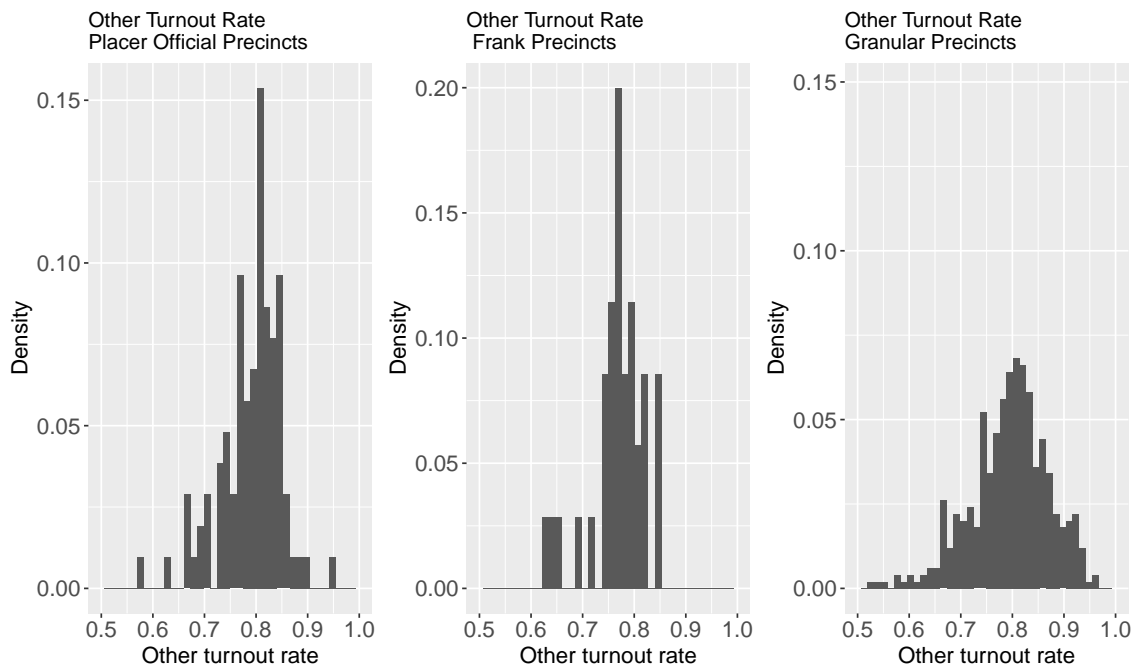


Figure 3: There is considerable variation in the turnout of individuals who do not register as a Republican or Democrat across Placer County precincts, regardless of how those precincts are defined.

## Frank's Analysis Merely Correlates a Variable With Itself

In Frank (2023) he describes additional age-based analyses, claiming that “Stated simply, *if one knows how many people are registered to vote in a particular precinct, one can very accurately predict the number of voters of every age in that precinct*” (Frank 2023, 3) (emphasis original). To make the prediction, Frank formulates a Placer County wide predicted turnout for each age. In a specific precinct, Frank then uses that predicted age-turnout rate and multiplies it by the number of individuals who are registered to vote and of a specific age in the precinct. He then correlates the precinct-level number of people of each age who he predicted would turn out to vote with the number of people of each age who actually turned out to vote. Frank asserts that “since the relationship between who is registered and who voted is so consistent across the entire county, one can use the sixth-order polynomial to predict the voter turnout by age in every precinct to a very high degree of accuracy” (Frank 2023, 3). Frank also makes similar claims about his ability to make predictions about turnout in counties within states. Frank calculates a state-wide predicted turnout rate for each age group. Then, in a specific county, he uses the statewide turnout rate and multiplies it by the number of registered voters of that age in that county. He then correlates his predicted turnout counts and the actual turnout counts. This corresponds to the analysis in Figure 8 in Frank (2023).

Frank, again, uses the wrong evidence to assess claims about a consistent turnout rate across precincts or counties. If Frank thinks the turnout *rates* across precincts in Placer County are the same, then he should assess whether each age group in each precinct in the county has the same turnout rate. And if he thinks the turnout *rates* across counties in a state are the same, then he should assess whether each age group in each county has the same turnout rate. Frank fails to produce this evidence for his Placer County analysis or for his analyses of counties across states.



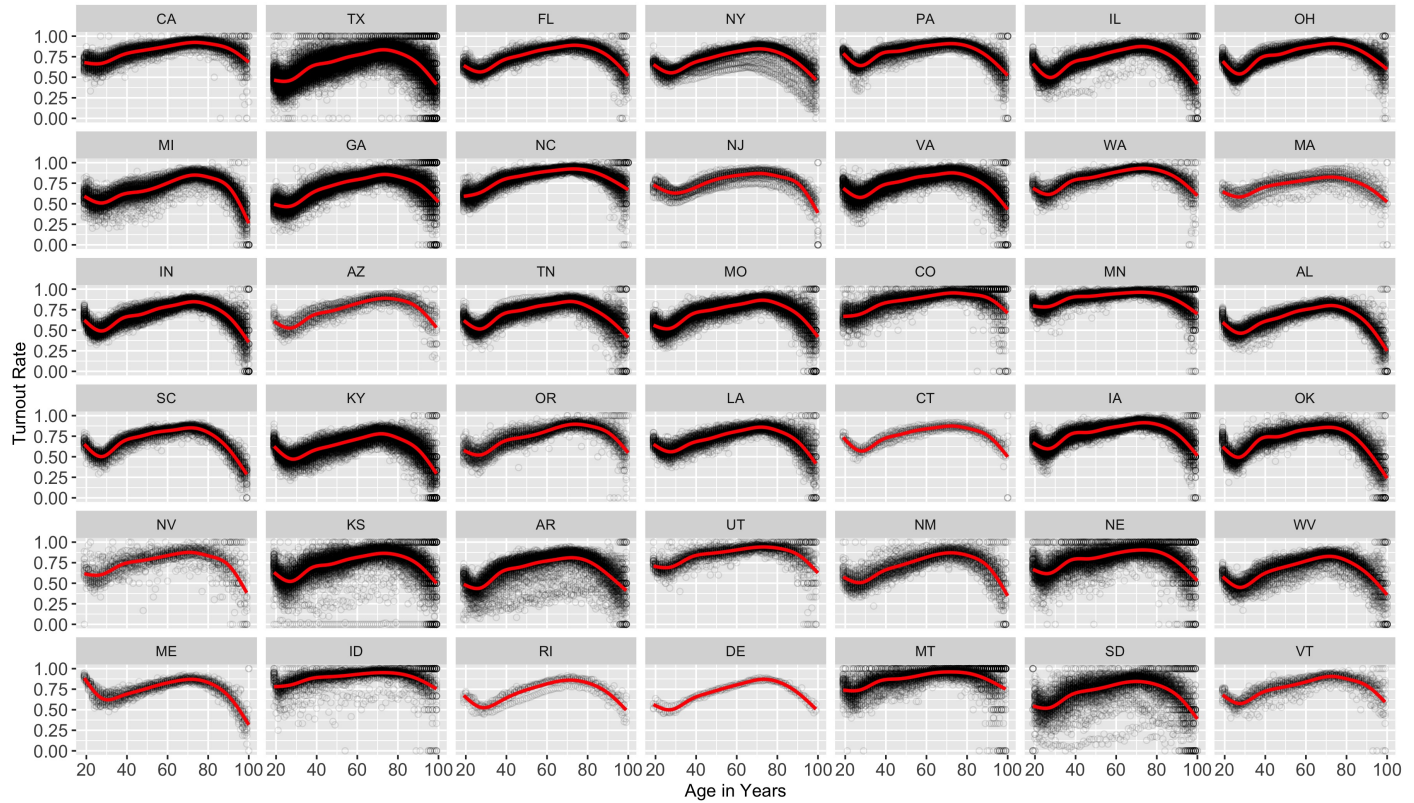


Figure 4: There is considerable county-to-county variation in the average turnout rate for each age group, contradicting Frank's claims that a single turnout rate can predict the turnout in each county.

Figure 4 shows that Franks' claims fail when appropriate evidence is used.<sup>2</sup> This plot shows the 42 states where I have age specific information from the voter file. Each point in the plot corresponds to the turnout rate for individuals of a particular age in a particular county. The red line is Frank's state-level prediction of the turnout rate. If Frank were correct, all the points in a state would align with this red line. But this is clearly false. There is considerable variation in the age specific turnout rates within a state.

Why, then, does Frank think he produces evidence with high correlation? It is because Frank continues to use incorrect statistical tests. Given Frank's predictions about turnout *rates*, a different test would be to calculate a correlation between the predicted turnout *rate* and the actual turnout *rate*. But this is not what Frank does. Instead, Frank uses the

<sup>2</sup>I focus on the county and state predictions because I do not have access to the age specific turnout for Placer County.

predicted turnout rate to make a prediction about the turnout *count* and compares that to the actual turnout *count*. We show in Grimmer et al. (2023) that using turnout counts inflates Frank’s estimated correlations. We provide extensive mathematical and simulation based evidence to demonstrate this and explain why focusing on counts inflates the correlation.

The core reason that Frank’s correlation of predicted and actual turnout counts causes an inflated correlation is that the correlation detects the unremarkable fact that age groups with more people registered to vote have more individuals who turnout to vote. When there is variation in the number of people of each age in a precinct or county, this inflates the correlation.

For intuition about how this inflation occurs, I performed a simulation and produced three simple plots.<sup>3</sup> In this example, I suppose that we’re assessing Frank’s claim that he can predict the turnout rate for each age in a precinct using the county-wide turnout rate for each age. To assess this claim, I examine a hypothetical precinct where I know that the correlation between the actual turnout rate and the predicted turnout rate, calculated using Doug Frank’s methodology, is 0.213. In the left-hand plot I examine a version of the hypothetical precinct where every age has almost exactly the same number of registered voters: there are 100 18 year olds, 100 19 year olds, etc, with the only exception being 101 40 year old residents. In other words, there is essentially no variation in the number of citizens of each age. On the horizontal axis I plot the actual number of voters from each age group and on the vertical axis I plot the predicted number from the simulation. Because there is essentially no variation in the size of the age groups in this version of the precinct, the correlation between the predicted turnout counts (using Doug Frank’s methodology) and the actual turnout counts is essentially equal to the correlation between the predicted and actual turnout rates.

But as we increase the variation in the number of individuals in each age group, the correlation between the predicted and actual turnout counts increases, *even though nothing*

---


<sup>3</sup>An expanded version of this simulation can be found at <https://electioninsights.org/learn>.

*else changes and the correlation between the predicted and actual turnout rates remain the same.* The middle plot is the same hypothetical precinct, but now I've increased the number of 40 year olds in the precinct. Of course, this increases the variation in the number of people in each age group. And as expected, the correlation between the predicted turnout count and the actual turnout count goes to 0.924. This occurs even though the correlation between the predicted and actual turnout rates remains unchanged at 0.213. It happens because the correlation "gives credit" to the prediction for correctly guessing that many more 40 year olds would turn out to vote, even though the correct guesses happened mechanically as the result of that group being larger.

And as we introduce much more variation in the size of age groups, we obtain correlations similar to what Frank calls "anomalous". The last panel shows the same precinct, but now I've introduced even more variation in age group size by making the number of 40 year olds even larger. Now I find a correlation between the predicted and actual turnout counts close to 1 (0.999), even though the correlation between the rates remains at 0.213. Again, this happens because the correlation between the predicted and actual counts is giving the correlation "credit" for knowing that there are more 40 year olds who are both registered to vote and then who turnout to vote. Of course, we could complicate this simulation in many ways but it will reveal the same basic principle: because there is variation in the number of people in age group, Frank's test merely uncovers the unsurprising fact that age groups with more registered voters have more voters who turnout.

This is far from a theoretical exercise. In fact, we demonstrate on [electioninsights.org](https://electioninsights.org) that in every county, we find that the correlation between the predicted and actual turnout count is larger than the correlation between the predicted and actual turnout rate. Not surprisingly, in Grimmer et al. (2023), we show that across states the average county-level correlation between predicted and actual turnout counts is higher than the correlation between the predicted and actual turnout rates in every state where we have data.

The reason for the artificially large correlations between the predicted and actual turnout



counts is that the *number of individuals registered to vote* appear in both sides of the correlation. Because there tends to be much more variation in the *number* of individuals who are of each age in a location than in their *turnout rates*, the *number* of individuals will increase this correlation. As a result, Frank's test is effectively correlating a variable with itself—the number of individuals who are registered to vote from each age group.

One way Frank claims to demonstrate the validity of his predictions is by asserting that the state-level turnout rate predictions only perform well within a particular state. He fails to provide evidence for this claim. To assess this claim, I used Frank's methodology, but for each county I calculated the predicted counts using the predicted turnout rates from every state, not just the county's home state. I then calculated the correlations between these out of state predictions and the actual turnout rates.

The results of this analysis are presented in the left-hand half of Figure 5. The top left-hand facet presents the distribution of correlations between the predicted and actual turnout counts using out of state turnout rates to make the prediction and the bottom left-hand facet

shows the distribution of correlations using the turnout rates from the county's same state to make the predictions. Both correlations are close to 1. In fact, the average correlation using the turnout rate from the county's same state is 0.99, while the average correlation using out of state predicted turnout rates is 0.984. In 16.9% of instances the predictions based on out of state turnout rates yields a higher correlation than the predictions based on the within state turnout rates. The right-hand plot correlates the county-level turnout rates with the state-level predictions about turnout rates. The top right-hand facet performs this correlation across state lines and the bottom right-hand facet performs this correlation in the same state. The correlations between predicted and actual turnout rates is much lower than the correlation between predicted and actual turnout counts, both within the same state and across state lines. This is because focusing on counts, rather than rates, artificially inflates Frank's correlations.

## **The “Bathtub” Analogy is Nonsensical and Fails to Meaningfully Address Any of My Critiques**

Frank argues that the critique of his methodology is mistaken, because it “misses the entire point, which is that *every county behaves unnaturally and identically*” (Frank 2023, 11) (emphasis original). Frank then defends his analysis with an analogy. The analogy involves sampling ping pong balls from a bathtub into bowls. Frank claims that because he finds a single number in a bowl in his analogy, he can conclude the sampling in his metaphorical bathtub is not random.

In no way does this analogy address the fact that Frank's analysis merely uncovers the unsurprising fact that groups with more people in them tend to have more people turnout to vote nor does it demonstrate that Frank's analysis reveals anomalous patterns in elections. Frank is assuming the conclusion: that he has identified an anomalous election pattern. But asserting an anomaly in this analogy is hardly evidence that “every county behaves

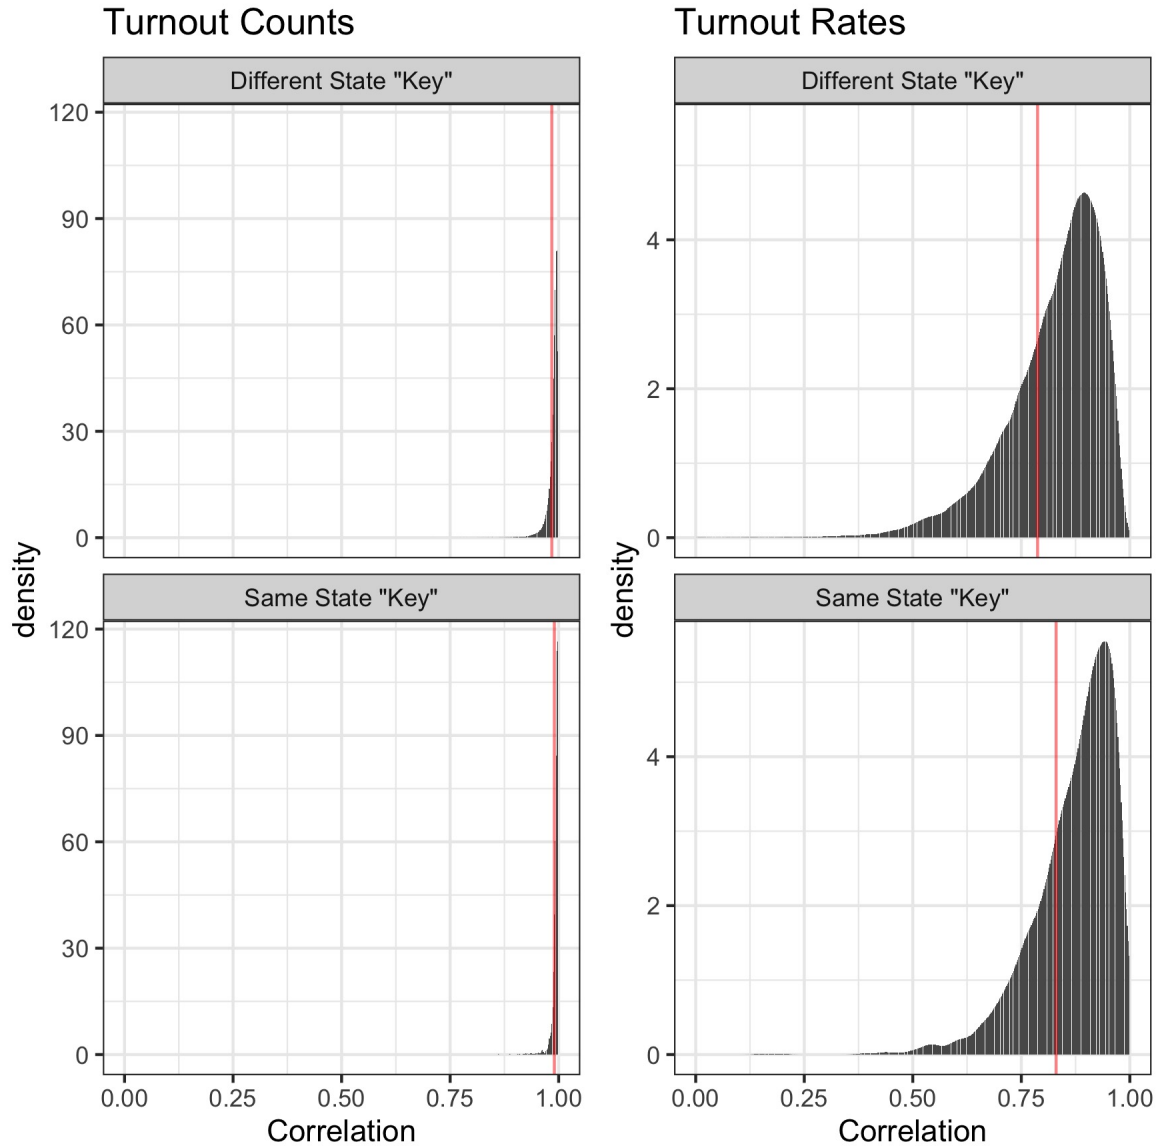


Figure 5: Average correlations between county-level predicted and actual turnout counts are very similar (left-hand plots), regardless of whether the turnout rates from the same state or different states is used to make predictions. This is the case even though the correlation between predicted and actual turnout rates are much lower, particularly when correlating predicted turnout rates across state lines.

unnaturally and identically.” In fact, our evidence shows that the patterns Frank categorizes as anomalous are quite normal and the result of Frank’s flawed interpretations of statistical analyses.<sup>4</sup>

## **Douglas Frank Analysis Fails As a Method to Identify Election Anomalies**

Douglas Frank’s analyses do not identify election fraud. This is because his methodology is based on a false premise and he fails to demonstrate that his methodology can accurately identify fraud in elections. I have shown that Frank’s analysis merely uncovers the unsurprising fact that groups with more people tend to have higher turnout rates. While Frank attempts to respond to the arguments, his response uses inappropriate evidence and is based on a nonsensical analogy. Neither response addresses the serious methodological flaws I have uncovered in his work.

In short, Frank fails to provide any evidence of anomalous results in U.S. or Placer County elections.

---

<sup>4</sup>Methodologically, the analogy is misplaced. There is no statistical, machine learning, nor social science literature where an analogy like this would be an appropriate way to address a critique.

## References

- Frank, Douglas (2023, June). Anomalous election turnout in placer county, california: Demographic analysis of the 2020 general election.
- Grimmer, Justin (2022). No evidence placer county elections are manipulated. Available here: <https://www.dropbox.com/s/a54kwsifx2gnq8c/GrimmerPlacerCounty.pdf?dl=0>.
- Grimmer, Justin , Michael Herron, and Matthew Tyler (2023). Evaluating a new generation of expansive claims about vote manipulation. Available here: <https://www.dropbox.com/s/92lp1gmqw2pei8m/Expansive.pdf?dl=0>.
- Grimmer, Justin , Margaret E Roberts, and Brandon M Stewart (2022). *Text as data: A new framework for machine learning and the social sciences*. Princeton University Press.