

Note: This is an expanded version of an article written by the 4 listed authors. This is Benjamin Leinwand's alternate version of the article. It is longer and includes more whimsical explorations than the standard version. Though the other authors' work was used to write this version, it does not necessarily represent their views nor reflect their publishing standards.

Winning an Election, Not a Popularity Contest

By: Benjamin Leinwand, Puyao Ge, Vidyadhar Kulkarni, Richard Smith

In two of the last five U.S. Presidential elections, the winner received fewer popular votes than their opponent. In 2000, George W. Bush received 47.9% of the popular vote to Al Gore's 48.4%, and in 2016, Donald Trump received 46.1% to Hilary Clinton's 48.2%. The remaining votes went to third party candidates. The political implications of the disparity between popular vote share and electoral college outcomes have been discussed at length. The extremes of this phenomenon, however, have been examined less thoroughly. In both 2000 and 2016, the popular vote winner won by less than three percentage points, but would it be possible to lose the popular vote by, say, ten percentage points and still win the election? What about 20%? Whether you refer to this extreme outcome as "unfair" or "efficient," it leads to the same question: what is the smallest percentage of the popular vote a candidate can garner while still winning the election?

This seemingly simple question can result in different answers depending on your assumptions about the universe of eligible voters, third party vote share, and the percentage of truly "up-for-grabs" voters. The answer also hinges on the practically implausible scenario of one candidate winning several states by exactly one vote, and losing several others by the maximum possible margin. Even though these outcomes are unrealistic, probing the boundaries can help us understand whether the most extreme possible outcomes of the current system are "par for the course" for any electoral college. We can contextualize the extreme scenario by grafting an electoral college onto the European Union, constructing an electoral college from metrics other than population, allowing for "fractional" electoral votes, and constructing the "fairest" and "least fair" possible extremal maps.

First, let's introduce how the electoral college works. While the US has 50 states, there are 51 territories of interest, as the 50 states and Washington, D.C. all have electoral votes which total to 538. Each territory gets a minimum of three electoral votes, with the rest of the 385 electoral votes assigned on the basis of population (D.C. cannot have more votes than any state – this caveat will come up later). Eight territories have the minimal three electoral votes, while California has 55. 49 territories¹ are "winner-take-all," that is to say the winner of the popular vote in the territory gets all of the territory's electoral votes. Maine and Nebraska assign one electoral vote to the winner of each of their districts (2 in Maine, 3 in Nebraska), and their remaining two electoral votes go to the statewide popular vote winner. The candidate who wins at least 270 electoral votes becomes the next president. By securing

¹ Maine and Nebraska assign one electoral vote to the winner of each of their districts (two in Maine, three in Nebraska), and their remaining two electoral votes go to the statewide popular vote winner.

269 electoral votes, a candidate can guarantee that no one else can win the election outright, at least preserving their own chance to win through a special process. We include percentages required to get to 269 electoral votes in parentheses to account for this possibility.

Since each state starts with three electoral votes regardless of population, and all states have an integer number of electoral votes, this gives rise to certain electoral votes representing fewer voters than others. We refer to the ratio $\frac{\text{Electoral Votes}}{\text{Popular votes}}$ as the “efficiency” of a territory (or candidate). To leverage the disparity in efficiencies among states, one can construct a simple heuristic, which goes roughly as follows: give the popular vote winner 100% of the votes in the least efficient states until they get to 268 or fewer electoral votes. Also give the popular vote winner 50% of the votes minus one in every remaining (more efficient) state. While this heuristic explains the general strategy for constructing an extreme scenario, there may be some electoral votes left over. For example, if taking the least efficient states in order results in a jump from 259 to 273 electoral votes, one can still include a state with seven electoral votes that is less efficient than the state we couldn’t add with 14. Extending this, perhaps we could remove a couple of the last added states to include slightly less efficient states, but get closer to 268 electoral votes, resulting in more total popular votes. The simple heuristic can’t account for these issues, but fortunately, we can use slightly more sophisticated methods like a Monte Carlo Algorithm and Integer Programs (see inset boxes).

We can start by comparing results depending on our definition of “population.” One choice is the “voting eligible population” (VEP) of each state, consisting of citizens over 18, but depending on the state, this can also disqualify felons and incapacitated people. Estimates of these quantities as of the 2020 primaries are available at the state level², but not available at the district level, so for Maine and Nebraska we have opted to split the voters evenly across districts. More narrowly, we can use the number of votes actually cast in the 2016 election in each territory as our reference for available popular votes. For the moment, we will assume that there are only two candidates and that all voters are “up-for-grabs.” We will discuss more “realistic” extreme scenarios at the end of this article. Using VEP, a candidate could win the election (get 270 or more electoral votes) with only 22.2% of the popular vote (22.1% if the candidate is

MONTE CARLO ALGORITHM

A simple algorithm can be used to get a “good” solution, but can’t guarantee an optimal one. Let S be the proposed set of territories won by the election winning candidate, and let N be a large number, like 1 million.

1. Let S = all 51 territories.
2. Randomly choose integer $K > 0$.
3. Randomly select K of the 51 territories.
4. Let S -new consist of S “swapping” all territories selected in step 3 (those that were in S are removed; those that were not in S are added to S).
5. If S -new has either (a) a larger population than S , or (b) fewer than 270 electoral votes, reject the change and keep S . Otherwise, replace S by S -new.
6. Go back to step 2 and run steps 2 - 6 for N repetitions.

You can also run this algorithm several times and choose the best available result.

² McDonald, Michael P. 2020. “2020 Presidential Nomination Contest Turnout Rates” *United States Elections Project*. October 1, 2020. <http://www.electproject.org/2020p>

INTEGER PROGRAM

An Integer Program (IP) is a constrained optimization problem where the decision variables are constrained to take integer values. An IP solver can provide a certificate that it has found an optimal solution, which is an improvement over the Monte Carlo Algorithm, but finding an optimal solution may take an extremely long time. Since there are only 51 territories, it is not computationally intractable to find optimal solutions to our problems using optimization software like CPLEX or SAS.

We can solve our initial problem with the following IP, which includes the “populations” and electoral votes in each territory as inputs.

Let: $1 \leq i \leq 51$, representing each territory.

$x(i) = 1$ if the winner gets a majority of popular votes in territory i , 0 otherwise.

$v(i)$ = Popular votes in territory i .

$E(i)$ = Electoral votes in territory i .

We solve the following problem:

Minimize:

$$(x(1)v(1) + \dots + x(51)v(51))/2$$

Subject to:

$$x(1)E(1) + \dots + x(51)E(51) \geq 270$$

$$\text{and } x(i) = 0 \text{ or } 1, 1 \leq i \leq 51$$

Even the more complex problems we discuss do not require much more sophisticated IP formulations.

only required to get 269 votes). If we use 2016 votes, a candidate can win with just 21.3% (21.2%)! Both the IP and the Monte Carlo produce the same solutions³.

The solution using 2016 vote has some peculiar features beyond the assumptions about voting patterns that were used to derive it. The inclusion of California and Texas in the most efficient set of territories contradicts the common notion that the way to efficiently amass electoral votes is to win all the small territories, and it seems unlikely that these two largest states would be won by the same party, since for decades California has been regarded as a Democratic stronghold with Texas just as strong on the Republican side. But perhaps that is part of the explanation why they feature in our optimal solution: both states have reduced turnout because they are regarded as a foregone conclusion, but that makes them more likely to be selected by our algorithm, which favors states with a low ratio of popular votes to electoral college votes.

Constructing the Most Extreme Scenarios:

At first blush, this seems like an absurdly small vote share for an electoral college winner. Perhaps the difference in efficiencies across states gives rise to such extreme values. One attempt to remedy this issue could be to keep the electoral college and let all states be winner take all (as it's not clear how to reallocate the districts in Maine and Nebraska), but reallocate the electoral

³ Using VEP, the set of territories worth at least 270 combined electoral votes but giving the minimum fraction of popular votes consists of: Alabama, Alaska, Arizona, Arkansas, Connecticut, Delaware, District of Columbia, Hawaii, Idaho, Illinois, Indiana, Iowa, Kansas, Kentucky, Louisiana, Maine, Maryland, Massachusetts, Minnesota, Mississippi, Montana, Nebraska, Nevada, New Hampshire, New Jersey, New Mexico, North Dakota, Oklahoma, Oregon, Rhode Island, South Carolina, South Dakota, Tennessee, Utah, Vermont, Washington, West Virginia, Wisconsin, and Wyoming. Using 2016 votes, add California, Georgia, and Texas to that set and remove Alabama, Illinois, Iowa, Maryland, Massachusetts, Minnesota, New Jersey, Oregon, Washington, and Wisconsin.

votes strictly by population. In this world, the number of electoral votes for each state is $538 \times \frac{\text{State population}}{\text{National Population}}$ and this value does not need to be an integer. This eliminates the 3 electoral vote minimum, makes all states equally efficient, and allows D.C. to have more electoral votes than certain states, but larger states will still have more votes than smaller states. In this case, the value we get is 24% (23.9%) for VEP, and 22.2% (22.1%) using 2016 votes.

One may now begin to suspect that any scenario with an electoral college gives rise to extreme possibilities. To see this, imagine if the playing field has been leveled as much as possible by creating 538 individual states of equal population, each with one electoral vote. Assume there are two candidates and all voters are up-for-grabs. Even in this contrived situation, one can use the heuristic to see that a candidate can win the electoral vote with just over 25% of the popular vote, by winning 270 states by exactly one vote, and getting zero other votes in the 268 other states. If each of these states consist of exactly one voter, we're back to a popular vote, and a candidate needs more than 50% of the popular vote to win. As the population of this fictitious country grows, though, the minimal winning vote share shrinks closer and closer to 25%, so we can think of 25% as an upper bound on the smallest percentage of the popular vote share of an electoral college winning candidate. In that sense, an electoral vote winning candidate getting as much as 25% is in some ways the "fairest" extreme outcome we could hope for! By comparison, the values we calculated above don't appear quite as bad as they first seemed, considering that these are the most peculiar possible outcomes.

With an upper bound in hand, we can also look for a lower bound. The upper bound relies on splitting the nation into equally populous territories. The lower bound has a somewhat stranger composition. To get there, we need to take a quick detour into how electoral votes are allocated to each state, called the Huntington-Hill method. The process is actually based on representatives and senators from each state, but for simplicity, we will explain how to get these values for the electoral college.

Each territory starts with three electoral votes. To distribute the remaining 385 votes, create a 51×385 table where each row represents a territory, and each column heading is calculated by the formula $\sqrt{c * (c + 1)}$, where c represents the column number. To get the value in each cell of the table, take the population of the corresponding state and divide by the column heading. For each of the largest 385 values anywhere in the table, the territory in whose row the value resides gets an additional electoral vote.

Getting back to the lower bound, instead of splitting the population equally across all territories, we put a single person in 49 of the territories with three electoral votes each, and split the rest of the people between two gigantic territories. Using the Huntington-Hill method, allocate the remaining people so that the larger one will have 268 electoral votes, but is one person short of receiving 269 electoral votes. The rest of the people reside in the final territory, which must have 123 electoral votes. Additionally, pack all people who are ineligible to vote into the second largest state, so they count towards allotting electoral votes, but cannot contribute to the popular vote. The US reported a total population of 328,239,523 in 2019, but currently has an estimated 233,053,576 VEP, which leaves 95,185,947 people who are unable to vote. If 226,036,192 people who are all eligible to vote live in the largest state, and 102,203,282 people live in the second largest state, that leaves only 7,017,335 eligible voters in the second state, one voter in all other territories, and crucially, gives the largest state 268 total electoral votes. This allows the losing candidate to spend votes maximally inefficiently, receiving no

votes from all of the ultra-efficient “one popular vote for three electoral votes” territories. If the electoral college loser wins all votes in the largest state and just under half the votes in the second largest state, that candidate would receive about 98.5% of the vote, leaving only 1.5% of the vote for the election winner. This calculation relies on the discrepancy between states’ total populations, according to which electoral votes are allotted, and their voting eligible population, who can cast popular votes. If instead we ignore all non-eligible voters and reallocate electoral votes solely based on VEP, this same two state packing method would allow the winner to receive just under 15.6% (15.4%) of the popular vote.

For those eagle-eyed readers waiting for Chekhov’s caveat, there is one even sillier situation. If we packed the whole US population (minus 50 people – one for every state) into Washington, D.C., beyond the feat of almost matching President Trump’s inauguration crowd, by law, D.C. cannot have more electoral votes than any state, so the votes would be split evenly across all territories. In that case, a candidate could receive only 26 votes – about one of every nine million eligible voters – and win the electoral college. Politically engaged readers may consider that this scenario could lead to increased calls for D.C. statehood, or perhaps secession.

Alternate Scenarios:

These upper and lower bounds are mathematically interesting and show the narrow range of these extreme possibilities, but they require such manufactured scenarios as to not provide much insight about our system in practice. That is to say, though one can concoct situations to achieve these upper and lower bounds, how do the current results compare to other “real world” possibilities?

As a first example, we can use the Huntington-Hill method and Mixed Integer Programming to see how extreme the electoral college could look were it imposed upon the European Union. In that case, the electoral vote winner would need to get more than 23.2% (23.1%) of the vote. However, this isn’t quite a fair comparison. The EU only has 27 member states, so only 81 of the 538 electoral votes are assigned without regard for population, as opposed to the 153 that are assigned in the US, which allows for greater efficiency for US candidates. To account for this, we can create a duplicate for each EU member state to get to 54 members. In that case, one could win their electoral college with 20.6% (20.3%) of the vote. One could object to our steadfast fidelity to 538 votes, but these results indicate at the extremes, the US cannot give rise to much more lopsided results than the EU would if it had an electoral college.

What if we allocated electoral votes by something other than population? President Trump is apparently fond of looking at the 2016 electoral map’s sea of red districts. What if territories were allocated electoral votes by their area, rather than their population size? Using the total area of each territory in square miles to distribute electoral votes, under winner-take-all-rules, a candidate can win with a vote share of 22% (21.9%). We could also cut out Citizens United’s middleman and let each dollar of GDP from each state literally get a vote. Money talks, so why shouldn’t it vote? A candidate could win an election carrying only 21.5% (21.4%) of the country’s GDP. Though these examples are unusual, at least the voting “population” is identical to the metric for allocating electoral votes. Electoral votes are apportioned to states based on their populations in 2010, so between 6 years of population changes,

different age and citizenship compositions of each state (as well as different laws governing voting access), and differing voting rates, using actual voters from 2016 allows for more lopsided results!

Realistic Scenarios:

The results presented thus far may terrify some, but these fears can be dismissed as entirely unrealistic. After all, no candidate will ever get 100% of the popular vote in any state! Especially in this polarized era, there may be relatively few true swing voters, so this kind of extreme minority rule is mere mathematical scare-mongering with no bearing on reality. With that in mind, we can incorporate more information to ground our (still extremely improbable) results in a more realistic context. The first piece of information is the *smallest vote share* each major party has received in any recent presidential election. This percentage represents the “floor” of support for each party, consisting of “solid” Democrats and Republicans. All voters not “solidly” in a camp are considered swing voters, and they are “up-for-grabs,” meaning either candidate can win any number of these voters. In each state, one party may have a higher floor than the other, which is a departure from the original assumptions, where no candidate is guaranteed any votes. For example, looking at the elections from 1992 to 2016, the smallest vote share for Democrats in California is the 46% Bill Clinton received in 1992, while Donald Trump’s 31.6% in 2016 is the smallest vote share for Republicans in California over that period, leaving 22.4% of California voters up-for-grabs. In North Carolina, the lowest vote share since 1992 is 42.7% for Democrats and 43.4% for Republicans, both occurring 1992, meaning only 13.9% of Tar Heel voters are considered swing voters. Furthermore, since candidates have a floor of support in each state, that should eliminate truly extreme results as described above.

We can also incorporate the *largest third-party vote share* from the past several elections, or equivalently, the largest percentage of votes that went to neither of the two major parties. A greater third-party vote share allows a candidate to win a state with a smaller percentage of the popular vote. For example, in 2016, Donald Trump won Utah while only receiving 45.5% of the popular vote, against Hillary Clinton’s 27.5%. An independent candidate, Evan McMullin, received 21.5% of the vote, allowing Trump to win with a relatively small plurality. In 1912, Woodrow Wilson won Idaho’s four electoral votes with only 32.1% (!) of the vote, while Presidents Teddy Roosevelt and William Howard Taft received 24.1% and 31% respectively, and the Socialist Eugene Debs picked off another 11.3%. As this example shows, there can be more than one “third-party” candidate, and in a race with many candidates with their own dedicated constituencies, a candidate can carry a state with a very small percentage of its popular vote. In practice, this is relatively rare in the general election for president, but it comes up frequently in party primaries with large fields of candidates. Though not winner-take-all, Bernie Sanders won the popular vote at this year’s Iowa Democratic caucus while receiving only 26% of the vote. Imposing a cap (but no floor) on third-party vote shares in each state means a candidate can’t win a lot of states with a very small vote share by benefiting from an unreasonably strong third-party candidate.

As we move to these “realistic” scenarios, we also change the question a bit. Instead of asking “what is the smallest percentage of the popular vote a candidate can garner while still winning the election,” we now ask “what is the largest percentage of the popular vote by which a candidate can lose while still winning the election?” When there are only two candidates, these questions are identical. However, when third-parties are present, minimizing the popular vote share for a winning candidate may not be the same as maximizing the losing candidate’s popular vote edge over the winning

candidate. For example, consider a very inefficient state (in other words, a state the electoral college winner should *lose* if they wanted to minimize popular vote share) with 40% of voters committed to both parties and up to 20% third party vote share. For the purposes of minimizing the election winner's vote share, it doesn't matter whether the candidate loses 40% - 60% or 40% - 41% with 19% third party vote share. However, if we want to maximize the gap between the loser's vote share and the winner's, the candidate should lose 40% - 60%. When third-party vote share can be high, answering the former question could allow a winning candidate to lose a handful of votes in exchange for their opponent losing many votes to third-parties, thereby shrinking the gap between them. The modified question seems more in the spirit of our inquiry. The answer to this latter question depends on which party wins, but also on our beliefs about support floors and third-party vote share.

We use the most extreme statewide results from recent elections as a proxy for each party's support floor and maximal third-party vote share in each state, allowing for two definitions of "recent." In the first case, we include the five elections between 2000 and 2016, as these elections have had relatively small national third-party vote shares. For the second case, we consider all presidential elections since 1992. In 1992, Ross Perot received 18.9% of the popular vote (he ran again in 1996, but only drew 8.4%). Though later third-party candidates have not drawn as many votes, it's not implausible to consider that certain individuals could run on an alternative ticket and garner as much support as Perot did. The most extreme achievable results based on these different assumptions, as well as the choice of population, are presented in the table below.

		Since 1992		Since 2000	
		VEP	2016 Votes	VEP	2016 Votes
D Win		42.6 - 55.2 (42.6 - 55.3)	42.9 - 55.0 (42.8 - 55.0)	46.4 - 53.2 (46.5 - 53.3)	46.7 - 52.9 (46.8 - 53.0)
		40.6 - 57.8 (40.0 - 57.3)	39.4 - 57.2 (39.4 - 57.3)	44.8 - 55.2 (44.8 - 55.2)	44.6 - 55.4 (44.6 - 55.4)
R Win					

It's no surprise that going back to 1992 allows for more extreme results, as that year featured an unusually large vote share for a third-party candidate. Additionally, given two more elections, each party gets two more opportunities in each state to post its own lowest vote share. It also may not be surprising to see that, no matter the assumptions, Republicans can win the election with getting far fewer votes than Democrats could. The distributional characteristics of the electorate that have enabled the GOP to win two elections while losing the popular vote this century also feature in these calculations. Still, even with "conservative assumptions," Democrats can also win the election while losing the popular vote by more than six points.

Of course, these are not truly conservative assumptions, as these extreme results still rely on one candidate picking up 100% of swing voters in some states. Furthermore, if there are fewer swing voters due to diverging party coalitions, using data even dating back to 2000 might inflate the number of up-for-grabs voters. However, this exercise is not intended to be in any way predictive of what might happen in the 2020 election. Instead, the purpose has been akin to a stress test, stretching possibilities until we are able to examine the limits of the electoral college. Considering the conceivable outcomes of a system can unearth surprising results, and perhaps even change how we think about the system in question. However, as we've shown, popularity may not be the only determinant of success in the US. In

fact, we have presented numerous scenarios, with varying degrees of plausibility, that show how it is possible to win the US presidential election with a minority of the popular vote. Of course, one simple change in the rules would make all our scenarios not merely implausible, but actually impossible: change the law so that the winner of the national popular vote becomes the president.