Pack-Crack-Pack: Gerrymandering with Differential Turnout

L. Bouton ⓡ G. Genicot ⓡ M. Castanheira ⓡ A. Stashko *

June 2023

ABSTRACT

This paper studies the manipulation of electoral maps by political parties, known as *gerrymandering*. At the core of our analysis is the recognition that districts must have the same population size but only voters matter for electoral incentives. Using a novel model of gerrymandering that allows for heterogeneity in turnout rates, we show that parties adopt different gerrymandering strategies depending on the turnout rates of their supporters relative to those of their opponents. The broad pattern is to *pack-crack-pack* along the turnout dimension. That is, parties benefit from packing both supporters with a low turnout rate and opponents with a high turnout rate in some districts, while creating districts that mix supporters and opponents with intermediate turnout rates. This framework allows us to derive a number of empirical implications about the link between partisan support, turnout rates, and electoral maps. Using a novel empirical strategy that relies on the comparison of maps proposed by Democrats and Republicans during the 2020 redistricting cycle in the US, we then bring such empirical implications to the data and find support for them.

JEL Classification Numbers: D72
Keywords: Redistricting, Gerrymandering, Electoral Maps, Turnout.

* Genicot and Bouton: Georgetown University, CEPR, and NBER, Castanheira: FNRS, Université Libre de Bruxelles and CEPR; Stashko: Emory University. Genicot and Bouton acknowledge NSF grant SES 2242288. We are grateful to Luke Miller for his research assistance. We thank Kfir Eliaz, Adam Meirowitz, Debraj Ray, Alexander Wolitzky, and seminar participants at the California Institute of Technology, Emory University, George Washington University, Hebrew University, HEC Montreal, London School of Economics, New York University, Queen Mary University London, Tel Aviv University, University College London, and the University of Chicago for their comments. ⓡ indicates that author names are in random order.

1. Introduction

Gerrymandering refers to the political manipulation of electoral maps: the strategic redrawing of electoral districts by the incumbent party to its advantage.[1] Most countries tackled this issue by delegating the redrawing of districts to independent electoral commissions (Stephanopoulos 2013).[2] In contrast, partly due to enhanced data and software, gerrymandering is stronger than ever in the United States. Beyond producing "bizarrely-shaped" districts, that are not geographically compact or cohesive, gerrymandering appears to impact policy-making. Empirical studies find that gerrymandering affects the quality of political candidates (Stephanopoulos and Warshaw 2020), the ideological position of legislators (Jeong and Shenoy 2022; Caughey et al. 2017; Shotts 2003), roll-call voting behavior (Jones and Walsh 2018), the ideological slant of the policies passed (Caughey et al. 2017), and the allocation of public resources (Stashko 2020). This impact provides an unfair advantage to the party of the gerrymanderer. Perhaps as a consequence, gerrymandering is viewed as a major problem of the U.S. political system by a majority of Americans (APNORC Center for Public Affairs Research 2021).

This paper focuses on the strategic incentives of gerrymanderers. Without understanding those incentives and what an optimally gerrymandered map may look like, we cannot identify a gerrymandered map, develop measures of gerrymandering, evaluate the effects gerrymandering could have on electoral outcomes and policy-making, and assess potential regulations.

To study those incentives, we propose a novel model of gerrymandering and then bring some of its predictions to the data. We take into account a simple but important fact that has been mostly overlooked by the literature: not all inhabitants of a district vote. Indeed, only 72% of the U.S. population was eligible to vote in 2020,[3] and, even among eligible voters, turnout rates vary substantially by, e.g., age, education level, and race.[4] And while districts must be equal in population size (counting voters and non-voters), only voters matter for electoral outcomes.

[1] The term is a portmanteau between the last name of Massachusetts Governor Elbridge Gerry and the word "salamander," alluding to the bizarre shape of districts when in 1812 the Republican-controlled legislature redrew the state senate districts in Massachusetts.

[2] Exceptions include Hungary, Lebanon, and most notably, the U.S.

[3] The remaining 28% were ineligible due to various reasons, such as being under 18 years old, being convicted felons, or not being U.S. citizens.

[4] As shown in Table F.5 and Figure F.2 in Appendix F.5, there are substantial variations in voter eligibility and turnout rates, both within states, and across and within partisan groups.

To illustrate how heterogeneity in turnout rates may change the incentives of gerrymanderers, consider the following simplified example. Let us imagine a state that is populated with 120 Democrats and 120 Republicans. These individuals must be allocated to four equally sized districts. When all individuals vote, the gerrymanderer's best strategy is to "pack" half of their opponents into one district and "crack" the rest of the population in the other three districts. A Republican, for instance, draws one district with 60 Democrats and no Republicans, allowing the creation of three districts with 40 Republicans and 20 Democrats. The Republican gerrymanderer then wins 3 out of 4 districts, even though they have support from half the population.[5] A Democrat gerrymanderer would follow the same "pack-and-crack" strategy to win 3 districts of their own.

The situation becomes starkly different if voters turn out at different rates. Suppose that each Republican votes with probability 1 and each Democrat votes with probability 1/2. In this case, the Republican gerrymanderer can win all four districts with a "crack only" strategy. If they construct four uniform districts with 30 Republicans and 30 Democrats, then the expected Republican vote share is 2/3 in each district. This illustrates the idea that gerrymanderers should exploit a turnout advantage by mixing higher-turnout supporters with lower-turnout opponents.

What should the Democrat gerrymanderer do in this case? The original pack-and-crack strategy is no longer optimal due to the turnout disadvantage: the three districts with 40 Democrats and 20 Republicans result in ties. With such a turnout disadvantage, the best strategy becomes one of "pack-and-pack": create two districts composed only of Democratic supporters, and two districts with Republican supporters. In this way, Democrats win the former two districts. This illustrates the idea that gerrymanderers should protect their lower-turnout supporters from higher-turnout opponents; a result that will be central to our analysis.

To analyze this problem in a substantially more general setting, our paper proposes a new model of gerrymandering that builds on a standard probabilistic voting model (see *e.g.* Lindbeck and Weibull 1987). The population is composed of various groups of individuals, which are characterized by their ideological lean—with democrats having a higher probability to prefer the democratic party than republicans do—and by their turnout rates (as in Strömberg 2004). Before the election, the gerrymanderer must draw the electoral map. This means allocating

[5] By contrast, drawing four districts with 30 Republicans and 30 Democrats offers a 50/50 chance of winning each district; that is 2 seats in expectation.

individuals from the different groups to districts, with the two-pronged constraint that each district should contain the same population mass and, naturally, that all of the population be allocated to a district. The gerrymanderer's objective is to maximize the expected seat share of the incumbent party in the state assembly.

Solving for the optimal map, we identify a "pack-crack-pack" strategy along the turnout dimension. The gerrymanderer's incentives are to isolate both their lower-turnout supporters and their higher-turnout opponents in packed districts and to mix intermediate-turnout supporters and opponents in cracked districts. Moreover, we find that the mixing of supporters and opponents in cracked districts should follow an assortative matching pattern: districts can be ordered along the turnout rates of both opponents and supporters. Finally, the model predicts a negative correlation between the average turnout rate of individuals of each partisan leaning and the probability that the gerrymanderer wins the district they are assigned to.

We then examine recent gerrymandering attempts to determine whether our model is supported by the data. To do so, we collect maps proposed by Republicans and Democrats in the most recent 2020 redistricting cycle. The district maps are combined with precinct-level data on population, turnout, and vote shares in recent presidential elections. This novel dataset, which includes the maps of 20 districting proposals over 10 states, allows us to compare redistricting proposals from both parties in the same state. The comparison of Democrat and Republican maps for a given state is key to attribute district characteristics to partisan strategy over other redistricting constraints and considerations.

We evaluate and find support for two fundamental theoretical predictions of the model. In Section 6.2 we test for the negative correlation between the turnout rate of the precincts composing a district and the anticipated vote share of the gerrymanderer's party in the district. In line with the theoretical prediction, we find that Democrat and Republican maps tend to treat the same precinct differently, depending on its turnout rate. In particular, a relatively low-turnout Democrat (Republican) precinct tends to get allocated to a district that leans more (less) strongly Democrat under the Democrat map than under the Republican map.

In Section 6.3, we test whether gerrymanderers exploit small modifications of district borders that increase their party's expected number of seats. To do so, we generate a large number of counterfactual districts by reallocating border precincts between pairs of adjacent districts. We then determine if a deviation is profitable by using expected vote shares to compute probabilities of winning districts as specified in our theoretical model. We compare the share of swaps

that are profitable for each party under Democrat and Republican proposals. We find that Democrats have fewer profitable deviations under their own proposal than under the Republican proposal. Likewise, Republicans have fewer profitable deviations under their proposal than under the Democrats'. Comparing within proposals, each proposal features a greater share of profitable deviations for the opponent's party than for the gerrymanderer's party. Finally, we find that heterogeneity in turnout rates plays an important role in the ability of our model to explain the patterns of profitable deviations, and therefore the locations of borders for Democrat and Republican proposals.

2. Literature

Theories of redistricting focus on partisan gerrymandering. Owen and Grofman (1988) is the germinal model of gerrymandering under uncertainty. Assuming two types of voters, opponents, and supporters, it rationalizes the so-called "pack and crack" strategy: to concentrate losses in as few districts as possible, the gerrymanderer should segregate ('pack') some opponents in districts that heavily favor the opponent. To win as many districts as possible, they should mix ('crack') the remaining opponents with supporters in districts that narrowly favor the gerrymanderer. In the presence of a high degree of uncertainty, however, even the mixed districts can be strongly in favor of the gerrymanderer. Gul and Pesendorfer (2010) and Friedman and Holden (2008, 2020) generalize this intuition in the presence of a continuum of ideological preferences and both aggregate and individual uncertainty.

Under a unifying framework, Kolotilin and Wolitzky (2020) find that pack-and-pair patterns —which generalize pack-and-crack—are typically optimal for the designer. When individual uncertainty dominates, the gerrymanderer packs the stronger opponents in some districts and pairs the others into cracked districts. When aggregate uncertainty dominates, a gerrymanderer may sort voters by the intensity of preferences. High-intensity Republicans are matched with high-intensity Democrats. This is known as matching extremes, or pairing, and is the districting strategy of focus in Friedman and Holden (2008, 2020). In all those models, individuals differ along one dimension, the intensity of their preferences for one party over the other, while individuals in our model differ along two dimensions, the intensity of preferences and the likelihood of voting.

Our model builds on that literature, with three important differences. First, as emphasized in the introduction, we extend the model to incorporate the turnout heterogeneity observed in

the electorate. Hence, in our model, individuals differ along two dimensions, the intensity of preferences and the likelihood of voting. Second, the model uses a standard probabilistic voting approach à la Persson and Tabellini (2000), which simplifies the analysis and connects our results to the broader political economics literature. Third, produces concrete implications that we bring to the data.

Introducing turnout heterogeneity proves qualitatively important: in contrast to our paper, none of the above-mentioned theories can rationalize the gerrymanderer packing their own supporters. Other potential explanations for the creation of safe districts in favor of the gerrymanderer's party include a high degree of aggregate uncertainty (Owen and Grofman 1988) and the need to protect incumbent seats. The latter explanation, known as incumbency gerrymandering, has found little support in recent empirical work.[6] Our paper adds one new explanation, which is that packing supporters can be optimal for the gerrymanderer in the presence of turnout disadvantages, even if there is very little aggregate uncertainty.

Another strand of the literature focuses on how to evaluate redistricting in practice. Most measures of partisan gerrymandering are designed to assess the extent to which a party packs and cracks voters.[7] These include the Efficiency Gap (Stephanopoulos and McGhee 2015), Mean-Median difference (McDonald and Best 2015), Declination (Warrington 2018), and Partisan Dislocation (DeFord et al. 2022). The Efficiency Gap is perhaps the most well-known among these. It measures the number of votes 'wasted' by one party relative to the other. The idea is that if one party packs and cracks, then the other party's votes are wasted because they win safe districts (votes in excess of 50% of the vote share are wasted) and lose competitive districts (all votes in lost districts are wasted). Note that our pack-crack-pack strategy would not necessarily be detected by the Efficiency Gap and other measures designed to detect pack and crack strategies. In particular, the creation of safe districts in favor of the gerrymanderer's party would conventionally be seen as 'wasteful' or non-strategic.

Though many measures of gerrymandering are based on the pack-and-crack strategy, few studies directly test whether gerrymanderers adhere to that strategy. Jeong and Shenoy (2022) offer some evidence suggesting that gerrymanderers indeed employ such a strategy. For instance, they observe that African Americans, who predominantly support Democrats, are allocated

[6] Gelman and King (1994) and Friedman and Holden (2009) find that, if anything, gerrymandering works against the interests of incumbents. Similar conclusions are reached by Abramowitz et al. (2006).

[7] Others evaluate the shapes of districts (Chambers and Miller 2010, Niemi et al. 1990), or compare outcomes to simulated districts (Chen and Rodden 2013, Gomberg et al. 2023.

differently across districts by Democrat and Republican gerrymanderers. Notably, they find that the packing of African Americans decreases when Republicans lose control of redistricting. This is in line with our finding that both the pack-and-crack strategy and the pack-crack-pack strategy are empirically relevant. Note that their empirical strategy also relies on the comparison of Democrat and Republican maps. However, they compare maps drawn during different redistricting cycles, following a shift in control of the redistricting process, whereas our analysis compares maps proposed within a single redistricting cycle.

Instead of detecting partisan strategies, other measures are designed to evaluate the normative properties of a map. Common measures include the bias, symmetry, and responsiveness of the seats-to-votes curve (Cox and Katz 2007, Grofman and King 2007, Katz et al. 2020). The seats-to-votes curve is the hypothetical relationship between the number of seats won by a party and the level of state-wide support. A map is biased if, for example, a party wins more than 50% of the seats whenever they win 50% of the votes. A map is asymmetric if parties could hypothetically win different seat shares given the same vote share. Last, a map has low responsiveness if the number of seats won by a party does not change much as the vote share changes. Katz et al. (2020) show that existing measures of partisan gerrymandering do not necessarily measure bias, symmetry, and responsiveness of a map. Importantly for our setting, they also highlight that both measurements of gerrymandering and estimates of a seats-to-votes curve are sensitive to the assumptions made about turnout rates. Thus, whether the intent is to detect a partisan gerrymander in practice or to determine how gerrymandering affects fairness in theory, one needs to understand redistricting strategies in the presence of heterogeneous turnout rates.

3. The Model

Two parties, the Democrat (D) and the Republican (R), will compete in a statewide election. The state is divided into J electoral districts, indexed by j. Each district elects one representative by first-past-the-post. Prior to the election, the incumbent has the opportunity to decide on the composition of each electoral district.

3.1. The Population

Each individual belongs to one of a finite number of groups, indexed by k. These groups differ along two dimensions: partisan lean and turnout rates.

Partisan lean is either Democratic or Republican: on average, Democrats assign a valence $\bar{\nu}^D \in \mathbb{R}$ to the Democratic party, which is higher than the valence $\bar{\nu}^R \in \mathbb{R}$ assigned by Republicans to the Democratic party: $\bar{\nu}^D > \bar{\nu}^R$. Within each group, individual voters will differ in their preference for either party; we return to this below.

The novelty of our model is to let population groups differ in turnout rates, denoted $\tau_k \in [0, 1]$. $N_k (> 0)$ denotes the population size of each of these groups. Naturally, population sizes N_k must add up to 1, the total size of the population.

To ease notation, we will often use subscripts d and d' to represent Democratic-leaning groups and subscripts r and r' to represent Republican-leaning groups.

3.2. Gerrymandering

The gerrymanderer designs an electoral map that allocates the population across electoral districts. An *electoral map* $\mathbf{n} := (n_{kj})_{\forall k,j}$ (with $n_{kj} \geq 0$) is a matrix that specifies how citizens of each group k are allocated to each district j.

We take the convention that the gerrymanderer belongs to party D. Their goal is to maximize the expected seat share of party D, $\pi(\mathbf{n})$, subject to the constraints that each district must have an equal share of the population $1/J$, and that all individuals are allocated to a district:

$$\max_{\mathbf{n}} \pi(\mathbf{n}) \text{ s.t. } \sum_k n_{kj} = \tfrac{1}{J},\ \forall j \quad \& \quad \sum_j n_{kj} = N_k,\ \forall k. \tag{1}$$

3.3. Probabilistic voting

We now turn to voting behavior, which underpins the probability of winning a district. In the tradition of probabilistic voting models à la Lindbeck and Weibull (1987), we assume that the gerrymanderer is uncertain about the voters' preferences. Uncertainty is only resolved at the time of the election.

Following (Persson and Tabellini, 2000, chapter 3), we consider two different types of shocks. First, an independently and identically distributed elector-level shock η_e toward party R captures preference heterogeneity among individuals within each group. Second, an aggregate shock δ captures the valence of party R over party D at the time of the election. The elector-level shock η_e follows a uniform distribution $U[-\frac{1}{2\phi}, \frac{1}{2\phi}]$. The aggregate shock δ follows another uniform distribution $\Gamma = U[-\frac{1}{2\gamma}, \frac{1}{2\gamma}]$.

Due to preference heterogeneity, not all Democratic leaning voters vote for party D, and not all Republican leaning voters for R. Conditional on turning out, an elector e with political lean $\bar{\nu}^P$ votes for D iff:

$$\bar{\nu}^P - \eta_e - \delta \geq 0. \tag{2}$$

Then, we have that the probability that party D wins district j is given by:[8]

$$\pi_j(\mathbf{n}_j) = \Gamma(\hat{\delta}_j) \text{ with } \hat{\delta}_j := \frac{t_j^D}{t_j^D + t_j^R}\Delta + \bar{\nu}^R. \tag{3}$$

where we denote by $t_j^D = \sum_d n_{dj}\tau_d$ the total turnout by Democrat-leaning individuals in district j, by $t_j^R = \sum_r n_{rj}\tau_r$ the equivalent total for Republican-leaning individuals, and by $\Delta = \bar{\nu}^D - \bar{\nu}^R$ the ideological gap between Democrat and Republican leaning voters. Since distributions are symmetric, the variable $\hat{\delta}_j$ represents the expected position of the median voter in district j.

In essence, (3) simply says that the gerrymanderer wins a district whenever its median voter prefers D to R. The expected lean in favor of D of the district's median voter is denoted by $\hat{\delta}_j$, and π_j is the probability that the aggregate shock δ remains below $\hat{\delta}_j$.[9] Importantly, $\hat{\delta}_j$ is increasing in the share of the district's overall turnout that leans Democrat, $t_j^D/(t_j^D + t_j^R)$, which we also refer to as the *effective share* of Democrats in district j.

In Appendix A, we introduce mild restrictions to guarantee that the resulting vote shares remain strictly between zero and one in each district, irrespective of their composition. Roughly speaking, this assumption requires that, relative to the ideological gap Δ, the density ϕ is sufficiently small, whereas γ is sufficiently large. Under this Assumption (A1) the probability that

[8]The probability that R wins district j is $1 - \pi_j$, making the problem symmetric for party R.

[9]This property is independent of the precise distribution of individual preferences. For instance, if individual preferences were normally distributed, the average –and hence median– district preference would remain the same weighted average $\hat{\delta}_j$.

party D wins district j boils down to:

$$\pi_j(\mathbf{n}_j) = \frac{1}{2} + \gamma\hat{\delta}_j.$$

The expected number of seats won by the Democrats being $\sum_j \pi_j$,[10] we can thus rewrite the gerrymanderer's problem (1) as:

$$\begin{aligned} \max_{\mathbf{n}} \pi(\mathbf{n}) = \sum_j \frac{\pi_j(\mathbf{n}_j)}{J} = \frac{1}{2} + \frac{\gamma}{J}\sum_j \hat{\delta}_j \\ \text{s.t. } \sum_k n_{kj} = \tfrac{1}{J},\ \forall j \quad \& \quad \sum_j n_{kj} = N_k,\ \forall k. \end{aligned} \tag{4}$$

Notice that the probability that Democrats win district j is linear in the effective share of Democrats in the district. For this reason, the gerrymanderer's objective boils down to maximizing a sum of Tullock contest success functions.[11]

3.4. Discussion of Key Assumptions

Before moving to the analysis of the gerrymanderer's behavior, we discuss five key assumptions of the model.

3.4.1 Objective

Our model assumes that the gerrymanderer draws the map with the sole objective of maximizing their expected seat share. This assumption aligns with the predominant objective assumed in existing gerrymandering models in the literature. Adhering to this approach thus ensures the comparability of our model with those of previous research. It is clear that there are various benefits for a party and its members of holding a larger number of seats and there is empirical evidence supporting this assumption (see, *e.g.* Jacobson and Kernell 1985; Incerti 2015; Snyder 1989 and see Genicot ⓡ al. 2021, pp. 3189-92, for a discussion).

[10] To see why, order the districts from most likely to elect a D candidate to least likely. For values of δ above $\hat{\delta}_1$, D has zero seats. For values of δ between $\hat{\delta}_1$ and $\hat{\delta}_2$, D has one seat (from district 1). For values of δ between $\hat{\delta}_2$ and $\hat{\delta}_3$, D has two seats (from districts 1 and 2), and so on. The expected number of seats is: $\sum_{j=1}^{J-1} j \times \left[\Gamma(\hat{\delta}_j) - \Gamma(\hat{\delta}_{j+1})\right] + J \times \Gamma(\hat{\delta}_J) = \sum_j \pi_j$.

[11] A contest success function is used by an increasingly large literature: see Tullock (1980); Hirshleifer (1989); Baron (1994); Skaperdas and Grofman (1995); Esteban and Ray (2001); Epstein and Nitzan (2006); Konrad (2007); Jia et al. (2013); Herrera et al. (2014); Bouton et al. (2018) among others.

Alternatively, we could assume that the gerrymanderer maximizes the probability of obtaining a majority of seats in the assembly like, for instance, Lizzeri and Persico (2001); Strömberg (2008). Such an objective would restore the incentive to draw a map following the typical "pack-and-crack" structure: creating a number of relatively strong districts that the gerrymanderer expects to win with the same probability and giving up on the others.

In practice, gerrymanderers' objectives go beyond pure partisan gains (Katz et al. 2020). For instance, parties may draw maps to protect incumbents or to hurt incumbents of the other party (and even encourage them to retire). While our model abstracts from those alternative objectives, our empirical strategy tries to take them into account. In particular, we focus on Congressional districts, which are drawn by state legislatures. Thus, no incumbents are directly involved in the redistricting process and there is no bonus for winning a majority of Congressional seats within a state.

3.4.2 Geographical and Legal Constraints

Our model, like much of the existing literature, places minimal constraints on the actions of gerrymanderers. Specifically, we abstract from legal constraints such as geographic contiguity, adherence to political boundaries, and compliance with the Voting Rights Act of 1965, which are typically observed to varying degrees in practice (Sherstyuk, 1998). These constraints will be revisited in the empirical section (see Section 6.1).

3.4.3 Distribution of the Aggregate Shocks

Our model assumes a uniform distribution for the aggregate shocks. This allows us to abstract from the "traditional" incentives to pack and crack one's opponents, and hence to isolate the novel incentives stemming from turnout rate differentials.

To see that heterogeneous turnout rates are critical when the aggregate shock is uniformly distributed, note that the gerrymanderer is indifferent between all electoral maps when turnout rates are equal. In particular, when population groups all share the same turnout rate, $\hat{\delta}_j$ is the average partisanship of the district. Under the assumption of a uniformly distributed aggregate shock, the expected probability of winning a district, $\pi_j = \Gamma(\hat{\delta}_j)$, is linear in $\hat{\delta}_j$. The combination of these two assumptions renders the objective tantamount to maximizing the average partisanship across all districts, which is invariant to the electoral map.

We show in Appendix E that assuming a (strictly) single-peaked distribution of the aggregate shock would reintroduce the traditional incentives to pack and crack one's opponents. With homogeneous turnout rates, our model would then be equivalent to the two-type model of Owen and Grofman (1988) or Kolotilin and Wolitzky (2020). With turnout heterogeneity instead, both the traditional pack and crack incentives and the novel forces highlighted in this paper coexist. We allow for a single-peaked distribution of the aggregate shock in the empirical analysis (Section 6.3).

3.4.4 Endogenous Turnout

Our model assumes that, within each group, turnout rates are exogenous and fixed. This captures the fact that a large part (if not most) of the heterogeneity in turnout rates depends on demographic and socioeconomic characteristics. For instance, in the 2020 U.S. Presidential election, only 72% of the population was eligible to vote (Census 2021, US Elections Project 2020). The remaining 28% were ineligible due to various reasons, such as being under 18 years old, being convicted felons, or not being U.S. citizens. Notably, there is significant cross-precinct heterogeneity in voting eligibility: at the 25th percentile of the precincts distribution, 63% of the total population are citizens of voting age. At the 75th percentile, this raises to 82%. Such variability is the primary reason for differences in voter turnout rates (defined here as the total number of votes divided by the total population) across precincts. Of the population that did not participate in the 2020 presidential election, 54% were ineligible to vote. Moreover, even among eligible voters, the propensity to vote differs across individuals for exogenous reasons (*e.g.* differences in education level, age, or ethnicity). And indeed, these determinants of turnout also vary greatly across precincts.

Yet, at the margin, the incentive to vote may also depend on the electoral map itself. In particular, turnout can be expected to increase when the electoral race is expected to be close in the district (see, *e.g.* Bursztyn et al. 2020 and Jones et al. 2023). In Appendix D, we show that our model can be extended in this direction without modifying the essence of our results.

3.4.5 Other Sources of Heterogeneity

We consider only one source of heterogeneity other than partisanship here – differential turnout. Other sources of heterogeneity would be similar, including swingness or information (Strömberg 2004), as long as they enter the sensitivity of the group (Genicot ⓡ Bouton ⓡ Castanheira

2021). Hence, the results that follow extend to a case where voters differ along a number of dimensions and where the gerrymanderer observes both the partisan lean and sensitivity of each population group.

4. Elementary Swaps

Since districts must be well apportioned and the number of districts is given, comparing two feasible electoral maps is equivalent to assessing the effect of a sequence of population exchanges or "swaps" across districts. The key question will be whether a swap (or a sequence of them) increases or decreases the expected number of seats controlled by the gerrymanderer's party. This section describes the basic properties of such swaps.

We denote by "$k \overset{i}{\rightleftharpoons}_j k'$ *swap*" the reallocation of some individuals of type $k \in \{d, r\}$ from district i to district j, in exchange for citizens of type $k' \in \{d', r'\}$ from district j.

Three types of swaps are possible: the reallocation of (i) two different types of Democrats ($d \overset{i}{\rightleftharpoons}_j d'$, or DD swaps for short), (ii) two different types of Republicans ($r \overset{i}{\rightleftharpoons}_j r'$, or RR swaps), and (iii) Republicans for Democrats ($d \overset{i}{\rightleftharpoons}_j r'$ or $r \overset{i}{\rightleftharpoons}_j d'$, or DR swaps). As a convention, we always denote by d and r the types who originate from district i and by d' and r' the types who originate from district j.

Property 1 (Swaps). *The effect of a $k \overset{i}{\rightleftharpoons}_j k'$ swap on the expected number of seats is:*

1. *concave for $d \overset{i}{\rightleftharpoons}_j d'$ swaps;*
2. *convex for $r \overset{i}{\rightleftharpoons}_j r'$ swaps;*
3. *concave (convex) in $d \overset{i}{\rightleftharpoons}_j r'$ swaps if $\tau_d > (<)\tau_{r'}$.*

The first two properties are very intuitive. Remember that, in the logic of a Tullock contest, the probability of winning a district is linear in the share of votes cast by d-types: $\frac{t^D}{t^D+t^R}$. That share is strictly concave in t^D and convex in t^R. Thus, whether the effect of a swap is concave or convex depends on how these swaps affect t^D or t^R. For instance, a DD swap involving individuals with different turnout rates has a positive and concave effect on t^D in one district and negative and concave effect on t^D in the other. The overall effect of such a swap is concave (formal developments are in Appendix B).

The effect of a DR swap is slightly more involved since it affects both t^D and t^R, but a similar logic applies. Take district j (the effects are symmetric in district i). Replacing Republicans

with Democrats in district j increases t_j^D, the numerator of the contest success function. The denominator, $t_j^D + t_j^R$, may increase or decrease, depending on the relative turnout rates of the Democrats and Republicans. If τ_d is higher (lower) than $\tau_{r'}$, then total turnout increases (decreases) in district j, and the swaps have diminishing (increasing) returns.[12]

When contemplating a concave effect, gerrymanderers will tend to perform swaps that reinforce their vote share in the districts where their vote share is lowest. Conversely, when the effect is convex, gerrymanderers will tend to perform swaps that reinforce their vote share in the districts where their vote share is highest. However, the profitability of a swap does not depend only on these concavity or convexity effects (Property 1). It also depends on the total turnout levels in each of the two districts.

To illustrate the effect of overall turnout, consider two districts with the same vote shares: $t_i^D/t_i = t_j^D/t_j$, but i features a lower turnout rate than j: $t_i < t_j$. The marginal effect of a given vote on the probability of winning the district is therefore higher in i than in j because a vote represents a larger share of total turnout. A gerrymanderer has an incentive to allocate (stronger) supporters to low-turnout districts, to amplify their effect. Conversely, allocating (stronger) opponents to high-turnout districts dilutes their negative effect on the expected number of seats won. In our example, the following swaps are thus profitable: a $d \overset{i}{\rightleftharpoons}_j d'$ swap with $\tau_d \leq \tau_d'$ (send relatively high turnout supporters to the low turnout district); a $r \overset{i}{\rightleftharpoons}_j r'$ swap with $\tau_r > \tau_r'$ (send relatively high turnout opponents to the high turnout district), and a $r \overset{i}{\rightleftharpoons}_j d'$ swap independently of the values of $\tau_{d'}$ and τ_r (send opponents to the high turnout district in exchange for supporters).

The concave/convex force and the overall turnout force are precisely the ones underpinning the numerical examples in the introduction: with just one d- and one r-group, only DR swaps are relevant. When $\tau_d > \tau_r$, swaps have a concave effect by Property 1. Starting with $t_i^D/t_i < t_j^D/t_j$, we must have that $t_i < t_j$, and hence both the concave/convex and the turnout-level effects invite $r \overset{i}{\rightleftharpoons}_j d'$ swaps. These incentives persist until the two districts are identical: the gerrymanderer wants to fully homogenize the districts. Conversely, when $\tau_d < \tau_r$, swaps have a convex effect. Starting with $t_i^D/t_i < t_j^D/t_j$, t_i must be larger than t_j. Thus, the gerrymanderer now wants to implement $d \overset{i}{\rightleftharpoons}_j r'$ swaps: the districts should be as segregated as possible. Here,

[12] The returns in one district are compounded by the same –diminishing or increasing– returns in the other district, since the swap experienced by i is the mirror image of that in j: t_i^D decreases, t_i^R increases, and total turnout in i moves in the opposite direction of that in j. Therefore, if the D-vote share displays a concave pattern in j, it must display a concave pattern also in i.

district i, should accumulate more r-types, and district j should accumulate more d-types, until at least one of these two districts is "packed", *i.e.* only contains d-types, or only r-types.

With two groups, these are comparatively simple situations, in which the concave/convex force and the turnout-level force are aligned. Yet, in other situations, the two forces may operate in opposite directions.

5. Optimal Gerrymandering

5.1. General Patterns

This section outlines general patterns characterizing an optimal map. For expositional clarity, we state our findings from the viewpoint of a Democrat gerrymanderer. By symmetry, equivalent results hold true for a Republican gerrymanderer.

If a district comprises solely supporters or opponents, it will be termed a "packed" district. A packed D district is defined as one where $\sum_d n_{dj} = 1/J$, whereas a packed R district is one where $\sum_r n_{rj} = 1/J$. Conversely, if a district includes both supporters and opponents, it will be called a "cracked" district.

Proposition 1 identifies a key criterion that guides the formation of packed and cracked districts: (i) the gerrymanderer never creates two districts that combine relatively low-turnout supporters and relatively high-turnout opponents, and (ii) the gerrymanderer never packs relatively high-turnout supporters in one district and relatively low-turnout opponents in another district. This generalizes the simple cases presented at the end of the last section.

Proposition 1. *No optimal map can have :*

(a) *two cracked districts i and j such that $\max\{\tau_d, \tau_{d'}\} < \min\{\tau_r, \tau_{r'}\}$ for any d, r types in district i and d', r' types in district j;*
(b) *a packed D district i and a packed R district j such that $\tau_d > \tau_{r'}$ for any d type in district i and r' type in district j.*

A common force drives the two parts of Proposition 1: gerrymanderers exploit to their advantage the differences in turnout rates between supporters and opponents. The election is more easily lost when a population of supporters gets outnumbered by an equally sized population

of higher-turnout opponents. Conversely, it is more easily won when the same population of supporters outnumbers an equally sized population of lower-turnout opponents.

Proposition 1(a) highlights a first consequence of that force: the gerrymanderer avoids mixing lower turnout supporters with higher turnout opponents. As already discussed, a key factor is the convexity of DR swaps when $\tau_r > \tau_d$. The open question was how convexity interacts with district-level turnout when there are multiple turnout rates. The proof of Proposition 1(a) shows that for any two districts that contain high turnout opponents and low turnout supporters, at least one feasible swap in the direction of more packing must be profitable. In particular, if the gerrymanderer has exhausted all profitable DD swaps and all profitable RR swaps, then there is a profitable DR or RD swap.

Proposition 1(b) reveals the flip side of the previous result: the gerrymanderer strictly prefers to mix high turnout partisans with lower turnout opponents. When Democrats have higher turnout rates than Republicans, by Property 1.3, the effect of a DR swap is concave, *and* the packed D district has a higher overall turnout. The concavity and the district-level turnout incentives are thus aligned to ensure that the gerrymanderer prefers to spread their supporters across both districts.

The next proposition addresses the question of how opponents of different types should be distributed across districts. It demonstrates that, when allocating opponents across two cracked districts, the gerrymanderer groups lower-turnout opponents in one district and higher-turnout opponents in the other.

Proposition 2. *In an optimal map, there cannot be two cracked districts, i with two types of republicans r and $\widetilde{r}$, and j with two types of republicans r' and $\widetilde{r}'$, such that $\tau_r > \tau_{r'}$ and $\tau_{\widetilde{r}} < \tau_{\widetilde{r}'}$.*

Proposition 2 follows from the convexity of RR swaps identified in Property 1. The gerrymanderer benefits from losing big in one district, concentrating high-turnout opponents there, in order to push its advantage in the other, now more favorable, district populated with lower turnout opponents. A special case of Proposition 2 is when there are only two types of republicans, r and r'. Then, the proposition implies that there cannot be two cracked districts where these two types coexist.

Next, Proposition 3 tells us that among cracked districts, there is positive assortative matching of supporters and opponents in terms of turnout rates: $\tau_r > \tau_{r'}$ implies $\tau_d > \tau_{d'}$ and vice versa.

Proposition 3. *In an optimal map, there cannot be two cracked districts i and j, respectively with types d, r and d', r' such that $\tau_d < \tau_{d'}$ and $\tau_r > \tau_{r'}$.*

From Proposition 1, we already know that, for any two cracked districts, supporters must have higher turnout rates than opponents. Proposition 3 adds that the gerrymanderer cannot benefit from combining their highest turnout supporters with the lowest turnout opponents. That is, we find that there are also decreasing returns to widening the $\tau_d - \tau_r$ gap in a district. Technically, the proof shows that after all DD and RR swaps are exhausted, if two districts still have negative assortative matching of opponents and supporters on turnout rates, then a DR swap that leads to positive assortative matching must be profitable.[13]

Additional pieces of notation are needed for the next proposition. We denote the lowest and highest turnout rates within each political faction as: $\underline{r} := \arg\min_r \tau_r$ and $\bar{r} := \arg\max_r \tau_r$, for the Republicans, and $\underline{d} := \arg\min_d \tau_d$ and $\bar{d} := \arg\max_d \tau_d$ for the Democrats.

The next proposition puts together the previous propositions to show that the optimal map must take a specific pattern, which we call '*pack-crack-pack*'. In particular, the gerrymanderer relies on a cutoff strategy, splitting the sets of supporters and of opponents according to their turnout rates.

Proposition 4. *In an optimal map, the allocation of voters to districts is characterized by two cutoffs, $\tau_{d^*} \in [\tau_{\underline{d}}, \tau_{\bar{d}}]$ and $\tau_{r^*} \in [\tau_{\underline{r}}, \tau_{\bar{r}}]$ such that:*

(1) *High-turnout supporters, with $\tau_d > \tau_{d^*}$, are assigned to mixed districts, whereas low-turnout supporters, with $\tau_d < \tau_{d^*}$, are assigned to packed districts;*
(2) *High-turnout opponents, with $\tau_r > \tau_{r^*}$, are assigned to packed districts, whereas low-turnout opponents, with $\tau_r < \tau_{r^*}$, are assigned to mixed districts.*

The intuition is simple. Supporters with higher turnout rates should be placed in districts where their votes matter the most, namely in districts where opponents are also present. Conversely, low-turnout supporters and high-turnout opponents are best assigned to packed districts, where this confrontation is absent. The proof proceeds as follows: start from a map that violates such a cutoff strategy, that is with a packed district that includes high-turnout supporters of type d

[13] This result complements the existing literature. The "segregate-pair" strategy identified in Kolotilin and Wolitzky (2020) also requires assortative matching but along the partisanship dimension: the most extreme opponents should be matched with the most extreme supporters. Note that they call this negative assortative matching since the voters who lean the most Democrats are then matched with the more Republican-leaning voters.

and a cracked district that includes low-turnout supporters of type d' ($\tau_d > \tau_{d'}$). A DD swap between those two districts is necessarily profitable since it does not affect the probability of winning the packed district (turnout rates are irrelevant in a packed district), whereas it does increase the probability of winning the cracked district, by widening the turnout gap between supporters and opponents.

Figure 1 illustrates the result for a case in which all three types of districts co-exist:[14] when the two cutoffs are strictly interior, there is a strictly positive number of districts packed with low-turnout supporters, cracked districts, and districts packed with high-turnout supporters. In line with Proposition 1(b), each cracked district mixes comparatively high-turnout supporters with comparatively low-turnout opponents. Moreover, following Proposition 3, this matching of supporters and opponents should be done in a positive assortative way.

This pack-crack-pack pattern complements the existing literature. First, it identifies differential turnout rates as a novel reason why gerrymanderers may prefer to pack and crack their opponents. In the classical "pack and crack" strategy proposed by Owen and Grofman (1988) (or its generalized "segregate-pair" version by Kolotilin and Wolitzky 2020), the gerrymanderer sorts opponents based on the intensity of their preferences, with the stauncher opponents allocated to packed districts. In contrast, in our model, the gerrymanderer sorts opponents based on their turnout rates, with the higher turnout rate opponents allocated to packed districts. Second, our model also identifies incentives for gerrymanderers to pack their lower-turnout supporters. Such a strategy cannot be optimal when individuals only differ along the partisan dimension. Indeed, packing supporters would conventionally be seen as an inefficient or 'wasteful' allocation of voters.

Finally, the following proposition identifies a negative correlation between turnout rates and winning probabilities in districts:

Proposition 5. *In an optimal map, all cracked districts can be ordered by increasing turnout rates and decreasing winning probabilities, so that for any two districts i and j with $i < j$:* $\max_{r \in i} \tau_r \leq \min_{r' \in j} \tau_{r'}$, $\max_{d \in i} \tau_d \leq \min_{d' \in j} \tau_{d'}$, *and* $\pi_i(\mathbf{n}_i) \geq \pi_j(\mathbf{n}_j)$.

[14] Section 5.2 below defines "robust" maps and provides sufficient conditions for the (co)existence of all three zones. Informally, and as displayed in the figure, $\tau_{\underline{d}}$ must be sufficiently low and $\tau_{\bar{r}}$ must be sufficiently high, with enough overlap between the turnout rates of supporters and opponents ($\tau_{\bar{d}} > \tau_{\underline{r}}$).

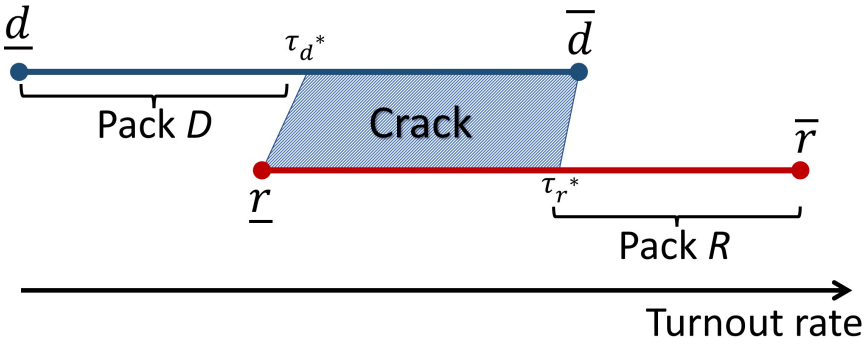


FIGURE 1. ILLUSTRATION OF PROPOSITION 5

The intuition for the proof is as follows: working by contradiction, consider two cracked districts i and j such that j contains both individuals with higher turnout rates (hence higher overall turnout) and a higher probability of winning. In that case, a DD swap is necessarily profitable. Indeed, by the concavity of these swaps, the gerrymanderer prefers to allocate higher-turnout supporters to the district with the lower probability of winning. The gerrymanderer also prefers to allocate higher-turnout supporters to the district with the lower overall turnout.[15] Hence, these two forces are aligned to render a DD swap profitable. To be optimal, the map must instead ensure that those two forces pull in opposite directions: the district with lower turnout rates must also be the district with a higher probability of winning.

It is worth noting that this ordering is also present across packed districts. Proposition 4 indicates that D-packed districts, which are most likely to be won by the Democratic party, must have lower turnout rates than Democratic voters allocated to mixed districts (in D-packed districts, $\tau_d \leq \tau_{d^*}$). Conversely, R-packed districts, which are least likely to be won by the Republican party, are composed of the highest Republican turnout rates ($\tau_r \geq \tau_{r^*}$).[16]

5.2. Robust Maps

To further describe the optimal map, we want to isolate the effects of the population structure, and abstract from the constraints imposed by a possibly small number of districts. Related

[15] Note that this implies that there is also a relationship between the district overall turnout and the district probability of winning. Yet, in the case of single-peaked aggregate shock, discussed in Appendix E, this district-level relationship does not hold as it also depends on the density of the aggregate shock.

[16] Remark that the gerrymanderer is indifferent to the specific distribution of types within D-packed districts and within R-packed districts. Creating identical districts or ordering them based on turnout rates would both be optimal.

models rule out such constraints by assuming a continuum of districts (Kolotilin and Wolitzky, 2020; Friedman and Holden, 2008, 2020; Gul and Pesendorfer, 2010).

Instead, we introduce the concepts of *map cloning* and *robust map* in Appendix C.1: first, let the gerrymanderer draw an optimal map given the existing number of districts. Second, clone this map any number of times by splitting each district into c identical clones of the original. This cloned map yields exactly the same payoff as the original map.

The question is whether the gerrymanderer can do strictly better after cloning. If the answer is positive, the map is considered not robust: what made it optimal was in part due to the restricted number of districts (see Appendix C.1 for a concrete example). If instead cloning does not open the door to further improvements, then the map is *robust*. We find that:

Proposition 6. *For a Democrat gerrymanderer:*

(1) *If* $\tau_{\underline{r}} < \tau_{\bar{d}}$, *a robust map must have at least one cracked district:* $\tau_{r^*} > \tau_{\underline{r}}$ *and* $\tau_{d^*} < \tau_{\bar{d}}$;
(2) *If* $\tau_{\bar{d}} < \tau_{\bar{r}}$, *a robust map must have at least one packed* R *district:* $\tau_{r^*} < \tau_{\bar{r}}$;
(3) *If* $\tau_{\underline{d}} < \tau_{\underline{r}}$, *a robust map must have at least one packed* D *district:* $\tau_{d^*} > \tau_{\underline{d}}$.

Concretely, the depiction in Figure 1 is a robust map if all three conditions in Proposition 6 are met. The left-most part of that figure may become empty if condition (3) is violated, etc. But, every time there is at least one group of supporters with a turnout rate higher than at least one group of opponents, a robust map will have a positive number of cracked districts.

6. Empirical Analysis

In this section, we examine recent instances of gerrymandering to assess the degree to which our model is substantiated by empirical data. Specifically, we evaluate two fundamental theoretical predictions of the model. Our first testable implication is whether we observe a negative correlation between turnout rates and expected vote shares across districts, as predicted by Proposition 5. We assess this reduced form prediction in Section 6.2. Second, the model predicts that gerrymanderers exploit profitable swaps, *i.e.* swaps that increase their party's expected number of seats. Section 6.3 utilizes the model even more closely to develop a test for this prediction.

For these empirical exercises, we collected data from 10 states in which both Republicans and Democrats submitted a redistricting proposal for congressional districts in 2020. Specifically,

we have 20 redistricting proposals from the states of Florida, Kansas, Louisiana, Maryland, Nebraska, Nevada, New Mexico, New York, North Carolina, and Pennsylvania, covering 113 congressional districts. We assemble precinct-level data to evaluate these redistricting proposals. Our final dataset includes 44,338 precincts across the 10 states (see Appendix F for more information).

6.1. Empirical Challenges

A primary challenge in testing the empirical support for our model lies in the need to reconcile it with reality. In our model, gerrymanderers allocate individuals to electoral districts. In practice, gerrymanderers are drawing maps using electoral precincts as the building blocks of districts.[17] The precincts are indeed the smallest possible geographic units at which electoral data are available. Hence, we shall take the precinct (instead of the individual) as our base unit for the empirical analysis.

A second challenge is to account for the geographic and legal redistricting constraints mentioned in Section 3.4.2 that the model abstracted from (e.g., geographic contiguity, respect for political boundaries, and adherence to the Voting Rights Act of 1965). Scholars have long noted that these factors make it difficult to identify a partisan gerrymander from an 'unintentional' gerrymander (Erikson, 1972; Chen and Rodden, 2013): Is a precinct in a particular district because this increases the seat share of the gerrymanderer or because of a legal or geographical constraint? For this reason, it is difficult to test if a gerrymanderer follows a pack-crack-pack strategy by simply looking at the composition of districts in a single map.

The key to overcome this challenge is to focus on *differences* across Democrat and Republican redistricting proposals. The idea is that if a precinct is assigned to a particular district due to legal constraints, both the Democrat and the Republican proposals will respect that constraint. Then, we observe no difference in that part of the state when comparing Democrat and Republican proposals. Any differences between the two proposals should be driven by partisan gerrymandering incentives, rather than by shared geographic or legal constraints.[18]

[17] Only 2% of the precincts in our sample are split across proposed congressional districts under either the Democrat or Republican proposals.

[18] As an example, consider the case of a majority-Latino community of interest. Under certain conditions, the Voting Rights Act would prohibit the community of interest from being divided across multiple districts. Majority-Latino precincts also tend to have relatively low turnout rates and to lean Democrat. The prediction from our model is for a Democrat gerrymanderer to pack the community of interest into a safe Democrat district. But if we observe this in practice, it could be due to the Voting Rights Act, rather than partisan gerrymandering, a false positive.

Ideally, we could compare an optimal Democrat proposal to an optimal Republican proposal for the same state at the same point in time. We approximate this by comparing Republican and Democrat proposals from the 2020 redistricting cycle.[19] These proposals were made during the 2020 redistricting process, either by party caucuses, partisan members of redistricting commissions, or by state legislators. We expect that the proposal by the majority party will be a strong partisan gerrymander, as it is likely to pass into law. The minority party in the state legislature also has incentives to draw a map in their favor. Their map can be used in negotiations if redistricting control is split across the two parties or as evidence in future lawsuits. Overall, the minority party wants to demonstrate that another map is feasible in which the majority party would win fewer seats. As a sanity check, we verify that the Democrat proposals are weakly better for the Democrat party than the Republican proposals in all states. Figure G.1 in the Appendix shows that in all the 10 states in our dataset, Democrats expect to win (weakly) more seats under the Democrat proposal than the Republican proposal (this comparison is strict in 8 of the 10 states).

A third challenge is that, for the sake of tractability and expositional clarity, our model assumes that the aggregate shock is distributed uniformly. As we discussed in Section 3.4.3, this assumption is meant to abstract from the traditional pack-and-crack incentives and allows us to focus the theoretical analysis on the new gerrymandering forces that turnout heterogeneity generates. Yet, one issue with that particular specification is that it predicts maps with counterfactual features. Most strikingly, it may be optimal for a gerrymanderer to draw a map with multiple cracked districts composed of a majority of opponents, that their party would lose almost certainly. This is an artifact of the uniform distribution, which implies that one more partisan voter in an almost surely lost district may have a bigger impact than in a close district,

Similarly, if the community of interest is in a strong Democrat district under the Republican proposal, we might improperly reject the theoretical prediction, a false negative. By comparing differences across proposals, we avoid both the false positive and the false negative in this example.

[19] Others have addressed this issue by comparing changes in redistricting proposals *across* time (e.g., Jeong and Shenoy 2022,Friedman and Holden 2009,Shotts 2003, Cox and Katz 2007). A limitation is that many other factors do change in a ten-year redistricting cycle, including the population and the number of districts per state. In 2020, 18 states had a different number of congressional districts than in the previous redistricting cycle. Given the small number of congressional districts per state, any change in the number of districts makes it difficult to compare maps over time. It is also rare to observe a change in partisan control within a state. In 2020, only two states (Arkansas and West Virginia) experienced a change in full partisan control over redistricting, and only Arkansas had both a change in partisan control and no change in the number of districts. Even for the 2010 cycle, when there was a nationally coordinated effort by Republicans to control the redistricting cycle, only 4 states switched control, and only 2 switched control with no change in the number of districts (party control of redistricting cycles by state and year are available at `https://redistricting.lls.edu/`).

by Property 1. As discussed in Appendix E, a straightforward way to remedy this is to consider an aggregate shock distributed according to a single-peaked distribution. This slight modification removes the incentive to create such weak districts and reintroduces traditional pack-and-crack incentives. Throughout our empirical analysis, we focus on this specification of our model, which combines traditional pack-and-crack forces and our new turnout-heterogeneity forces.

6.2. A Reduced-Form Test: Turnout and Vote Shares

Proposition 5 and the discussion that follows predict a negative relationship between the turnout rate of the precincts (the individuals in our model) composing a district and the probability that the gerrymanderer's party wins the district. Importantly, that relationship holds only for precincts of a given partisan lean.

Take all precincts of a given partisan lean in a state and order them by increasing turnout rates. Now, consider the probability that the Democrat wins the district to which each precinct is assigned. This probability of winning will then be decreasing with the turnout rates under the Democrat proposal, but increasing under the Republican proposal. Thus, the *difference* in the probability that Democrats win a district in the Democrat proposal versus the Republican proposal is decreasing in turnout rates, within Democrat- and Republican-leaning precincts. Given that the probability of winning a district is increasing in the anticipated vote share, we have the following empirical prediction.

Empirical Prediction 1. *Within precincts of a given partisan-lean in a given state, there is a negative relationship between the precinct turnout rate and the difference in the expected Democratic vote share of its assigned district under the Democrat versus Republican map.*

To help with the intuition of this prediction, consider a low turnout Democrat-leaning precinct. Democrats should put the precinct in a packed district (with low turnout) and Republicans should put it in a cracked district (together with higher turnout supporters). The difference in expected Democrat vote share across the two maps (Democrat proposal - Republican proposal) will be positive. Next, consider a high-turnout Democrat-leaning precinct. Democrats should assign it to a cracked district (together with lower turnout Republicans), while Republicans should assign it to a packed Democrat district (all with high turnout). The difference in expected Democrat vote share will be negative. Thus, the difference in expected Democrat vote

shares is decreasing in turnout rates, among Democrat-leaning precincts. The same exercise works for Republican-leaning precincts.

6.2.1 Empirical Methodology

To test the above prediction, we measure the turnout rate and partisan lean of each precinct and estimate the following empirical specification:

$$\textit{Difference in Democratic Vote Share}_{kps} = \beta\ \textit{Turnout}_{kps} + \theta_{ps} + \epsilon_{kps}, \tag{5}$$

where $\textit{Difference in Democratic Vote Share}_{kps}$ is the difference in the district-wide expected Democratic vote share under the Democrat proposal versus Republican proposal for a precinct k, of partisan-lean p, in state s. For Democratic vote share, we use the average votes of the two most recent presidential elections prior to redistricting, and aggregate votes to the district level.

For partisan-lean, we use party registration to classify a precinct as Democrat-leaning, Republican-leaning, or weakly partisan. We focus on precincts that lean clearly toward either the Democrat or Republican party and exclude weakly partisan precincts from the analysis. This is to keep in line with treating a precinct as a group with a known partisan leaning and turnout rate. We use the 25th and 75th percentiles of the two-party Democrat share of registered voters as thresholds. This implies a precinct is Democrat-leaning if the percent of Democrats is above 78% and is Republican-leaning if the percent of Democrats is below 38%.[20] We show that results are not sensitive to the thresholds used to define partisan-lean in Appendix G.2.

The variable $\textit{Turnout}_{kps}$ is one of the three measures of turnout defined in the Appendix F.4: past turnout, predicted turnout, and citizens voting age population (CVAP). We consider these three different measures to address concerns that turnout is potentially endogenous to redistricting. Predicted turnout is our preferred measure.[21] θ_{ps} is a vector of state-specific partisan-lean fixed effects, and ϵ_{kps} is an error term.

The coefficient of interest is β, which we predict is negative. We include state-specific partisan-lean fixed effects, θ_{ps}, so that the variation used to estimate β comes from comparisons within

[20] In Figures G.2 and G.3 in Appendix G.2 we show that there is a bimodal distribution of percent Democrats across all precincts, as well as within several states.

[21] We prefer predicted turnout over past turnout because we can more confidently rule out endogeneity concerns. We also prefer predicted turnout over percent CVAP since the latter is a crude proxy and gerrymanderers are likely aware of other salient correlates of turnout like race and income.

TABLE 1. Turnout and difference in district-level Democratic vote share.

	(1)	(2)	(3)
Percent CVAP	-0.006		
	(0.004)		
Predicted turnout rate		-0.021***	
		(0.006)	
Turnout rate			-0.012**
			(0.005)
N	21166	21166	21166
R^2	0.094	0.095	0.094
Outcome variable mean	-0.014	-0.014	-0.014
State-partisan lean FE	X	X	X

Note: OLS estimates. The dependent variable is the difference between the (two-party) Democratic vote share of the district under the Democrat proposal and under the Republican proposal (D map - R map). Percent CVAP is the percent of total population that are citizens above age 18. The turnout rate is the total number of votes for Democrats and Republicans divided by the total population. The measure of predicted turnout rate uses only demographic and socioeconomic variables to predict the turnout rate. $*$ p $<$ 0.1, ** p$<$ 0.05, *** p$<$ 0.01. Robust standard errors in parentheses.

Democrat-leaning and within Republican-leaning precincts in the same state, in line with the theory.

6.2.2 Results

Table 1 reports the coefficients from OLS estimations of Equation (5) for each of the three measures of turnout. In all regressions, the measure of turnout is negatively correlated with the difference in Democratic vote share, in line with the theoretical prediction. The correlations are stronger and more precise for predicted turnout and realized turnout than for percent CVAP. As mentioned above, our preferred measure of turnout is predicted turnout. The point estimate in column (2) suggests that a one standard deviation increase in predicted turnout (10 p.p.) is associated with a 0.2 p.p. decrease in the difference in the expected Democratic vote share of districts across proposals (which represents 14% of the mean difference). In Appendix G.2, we show that the negative correlation is robust to alternative definitions of partisan-lean. The coefficients for all three measures of turnout are negative if we include precincts with as little as 55% support for a given party, and up to 90% for a given party. As we restrict attention

to more strongly Democrat- and Republican-leaning precincts, the magnitude of the negative correlation increases. Overall, the negative correlations between turnout rates and differences in expected vote shares supports the idea that gerrymanderers exploit differences in turnout rates in a pattern consistent with the predictions of the theoretical model.

6.3. Swapping Precincts at the Border

In this section, our objective is to test whether, in practice, gerrymanderers fully exploit profitable deviations (*i.e.* , changes in the map that would increase their party's expected number of seats won). Testing this prediction requires taking the model seriously and creating counterfactual districts. However, due to the geographical and legal limitations discussed earlier, there will always remain some apparently unexploited profitable deviations. Consequently, we compare profitable deviations within a party across different proposed maps:

Empirical Prediction 2. *The share of unexploited profitable deviations for a party should be lower under its proposal than the one of the other party.*

Naturally, the number of possible counterfactual maps is extensive. Hence, we concentrate on counterfactual districts created by redistributing precincts along district boundaries. This has the added benefit of increasing the likelihood of feasibility of the counterfactual scenarios, in particular in view of the contiguity constraint.

Counterfactual Maps

For each pair of adjacent districts, we identify all precincts along the border, then independently randomly assign each precinct to one of the two districts with 50% probability. A random allocation of precincts is considered a feasible counterfactual map only if it satisfies the equal population constraint.[22] For a border with n precincts there are 2^n possible redistrictings. The median number of precincts per border is 20, but some borders have hundreds of precincts. Due to computational limitations, we sample until we generate 1,000 feasible maps per border, resulting in 353,000 simulated maps.[23]

[22] We impose strict population constraints for the counterfactual districts, as in reality. The total population of the counterfactual district must be within 1% of the size of an ideal district in the state. The size of an ideal district is the total population of the state divided by the number of districts. Given these tight population constraints, we reassign precincts one border at a time. Otherwise, assessing a large number of feasible counterfactual maps is a computational challenge.

[23] Moreover there are a few borders for which there is no feasible swap due to the strict population constraints. We exclude such borders from the analysis. Eventually, we simulated 179,000 maps for Democrat proposals and

Profitable Deviations

We can determine if a deviation is profitable by mapping vote shares into win probabilities. Let $\pi_1 + \pi_2$ be the sum of the probability of winning the two districts along a border under the baseline proposal and $\tilde{\pi}_1 + \tilde{\pi}_2$ be the same for the counterfactual map. A counterfactual redistricting is a profitable deviation if $\tilde{\pi}_1 + \tilde{\pi}_2 > \pi_1 + \pi_2$. As discussed above, in this empirical analysis, we focus on the specification of our model with single-peaked aggregate shocks. In particular, we consider a mean-zero normally-distributed aggregate shock. This creates an S-shaped π_j, consistent with the idea that gerrymanderers do not stand to gain much from improving their expected vote share in a district that is won or lost with near certainty. Their focus is mainly on more contestable districts.

To calculate the probabilities of winning, we need to choose a value for σ, the standard deviation of the aggregate shock. We calibrate the standard deviation of the aggregate shock using the state-wide standard deviation of the Democratic two-party vote share in presidential elections from 2008-2020,[24] the values range from 0.01 to 0.02 (see Appendix G.3). These values imply that gerrymanderers only meaningfully benefit from swaps in very competitive districts. To clarify that our results do not depend on the specific value, we depict the results for a range of values of σ between 0.005 and 0.3 in Appendix G.4. To understand the meaning of this range, first note that a district with 60% of the expected vote share is won with greater than 99% probability if $\sigma = 0.005$ and with probability 63% if sigma is 0.3.

6.3.1 Results

Figure 2 reports the share of swaps that are profitable for each party under Democrat and Republican proposals. There is a clear pattern: Democrats have fewer profitable deviations under their own proposal (30% of swaps) than under the Republican proposal (39% of swaps). Likewise, Republicans have fewer profitable deviations under their proposal (28%) than under the Democrat's (37%). Similarly, each proposal features a greater share of profitable deviations

174,000 for Republican proposals. The difference results from the fact that the number of borders and the number of precincts per border differs between the two maps.

[24] In our model, the standard deviation of the normal distribution is equal to the standard deviation in vote shares, divided by ϕ. Thus, we over-estimate σ if there is a relatively large degree of individual-level uncertainty and we under-estimate σ if there is a relatively low degree of individual-level uncertainty. Rather than calibrate ϕ, we present results over a range of σ. We use presidential elections rather than congressional elections to avoid imputing values for uncontested races. We use more presidential elections than in the precinct-level analysis since precinct borders change over time.

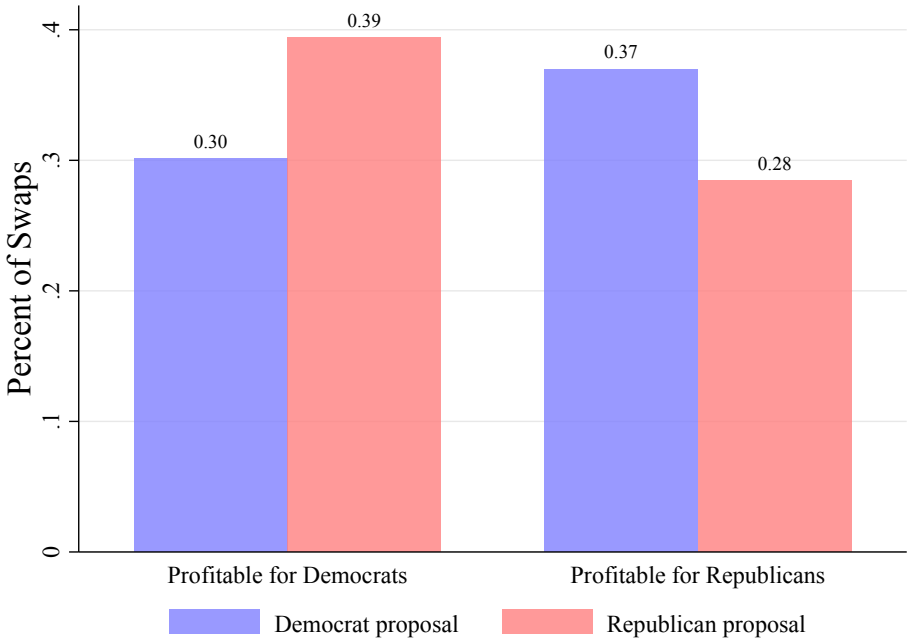


FIGURE 2. SHARE OF PROFITABLE SWAPS. The y-axis is the percent of feasible counterfactual maps ('swaps') that are profitable for Democrats (blue bars) and Republicans (red bars). A swap is profitable if the expected number of seats is higher than under the baseline proposal. To compute the expected number of seats we assume that the aggregate shock is normally distributed with mean zero and standard deviation calibrated using statewide presidential election returns from 2008-2020 (see Appendix G.3, values range from 0.01 to 0.02).

for the opponent's party than for the gerrymanderer's party (for instance, in Figure 2, the Democrat proposals in blue yield 37% profitable deviations for the Republicans, and 30% for the Democrats).

The effect of most swaps on the gerrymanderer's payoff (the sum of the probability of winning the two districts) is tiny.[25] The median profitable swap affects the payoff by 0.01 percentage points. In Figure G.5 in Appendix G.4 we show that the pattern in Figure 2 holds if we exclude swaps with negligible effects on payoffs.

[25]Some swaps have zero effect on the payoff, despite our measure being precise up to 24 decimal digits. This is why the bars in Figure 2 do not add up to 100%.

Finally, it is important to determine the role that heterogeneity in turnout rates plays in the ability of our model to explain the locations of borders on Democrat and Republican proposals. To do so, we repeat the swaps exercise, measuring the effect on the probability of winning the seats, but as if all individuals had the same turnout rate. Measuring what the vote share would be without turnout heterogeneity requires some assumptions. One approach is to assume that the vote share reflects the preferences of the total population. This is true only if there is no turnout rate heterogeneity across parties. In practice, we take the two-party vote share and multiply it by the total population to measure 'Scaled Votes' for Democrats and Republicans in each precinct. That is, moving a precinct with population n from a district i to another district j always translates in a transfer of n votes between these districts. In a second approach, we use party registration as a proxy for the partisan preferences of the population. We similarly multiply the share of registered voters that are Democrat or Republican by the total population to compute the 'Scaled Registered Voters' in each precinct. We then repeat the empirical swaps exercise, computing payoffs using the two scaled measures.

The purpose of this exercise is to measure the importance of turnout heterogeneity in assessing profitable deviations. Indeed, as detailed in Section 4, the profitability of a swap depends on turnout and partisan lean differences between districts. By artificially suppressing turnout heterogeneity, we focus on the incentive to exploit differences in partisan lean, *i.e.* , the standard pack-and-crack strategy. To the extent that turnout heterogeneity matters for gerrymanderers' payoffs in practice, we expect to find, at best, weak evidence in support of Empirical Prediction 2 when using scaled votes or scaled registered voters to measure payoffs.

Figure 3 shows that Democrats and Republicans each have *more* unexploited profitable swaps under their own proposal than under their opponent's proposal when using scaled voters to measure payoffs. For scaled registered voters, there is mixed evidence in support of prediction 2. Democrats have fewer profitable swaps under their own proposal than under the Republican proposal, but Republicans have more profitable swaps under their own proposal than under the Democrat proposal. Note also that, when using scaled turnout, the gerrymanderer's party has more profitable swaps than their opponent does in their own proposal. Overall, when we shut down heterogeneity in turnout, the patterns in Figure G.1 weaken (in the case of scaled registered voters) or reverse (in the case of scaled votes).[26] The reason is simple: it is difficult to determine if a swap is profitable without taking turnout into account. Of the swaps

[26]The patterns in Figure 3 are not sensitive to excluding negligible swaps or to varying the value of the standard deviation of the aggregate shock (Figures G.7 and G.8 in Appendix G.4).

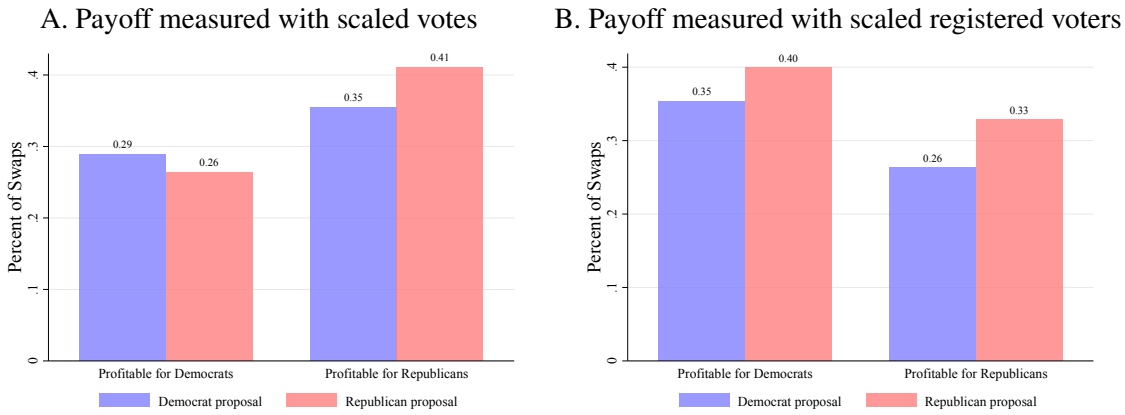


FIGURE 3. PROFITABLE DEVIATIONS WITHOUT TURNOUT DIFFERENTIAL. This figure compares profitable deviations when we measure $\hat{\delta}$ with scaled votes or scaled registered voters.

that are profitable for the gerrymanderer's party using scaled votes, only 57% are profitable using votes. Similarly, of the swaps that are profitable using scaled registered voters, 50% are profitable using votes.

7. Conclusions

This paper studies the strategic incentives of gerrymanderers when drawing electoral maps. Our approach introduces a novel aspect by considering the fact that not all individuals vote. We first present a formal theory of redistricting, taking into account heterogeneity in turnout. We show that a gerrymanderer aiming to maximize their party's number of seats allocates supporters and opponents quite differently across electoral districts. Specifically, low-turnout supporters are packed into safe districts to prevent their votes from being diluted by higher-turnout opponents, thus ensuring their influence on the final vote outcome. Conversely, high-turnout opponents are concentrated in disadvantaged districts, reducing their influence and concentrating losses in a smaller number of districts. Other population groups are allocated to mixed districts, deliberately associating opponents with higher-turnout supporters. We call this strategy "pack-crack-pack": going from low to high turnout rates, packing supporters, creating cracked districts that mix supporters and opponents, and finally packing very high-turnout opponents. Our model also predicts that districts can be ordered from lowest to highest turnout rates, with the probability of winning each district decreasing as average turnout increases.

In the second part of the paper, we empirically test two key predictions of the model: (i) a negative correlation between turnout rates and expected vote shares of the gerrymanderer's party across districts; (ii) gerrymanderers exploiting any swaps of individuals across districts that would increase their party's expected number of seats. To test these predictions, we compared the actual proposals put forth by both the Democrats and the Republics in ten U.S. states during the 2020 redistricting cycle. We found patterns in the data in line with both our predictions.

Turning to future research, we anticipate that these findings could contribute to refining the ex post measures of gerrymandering. As discussed in Section 2, most existing measures of partisan gerrymandering are designed to assess the degree to which a party packs and cracks voters. Yet, we have shown in this paper that, when there is heterogeneity in turnout rates, gerrymanderers may opt for a pack-crack-pack strategy instead of strictly adhering to the traditional approach. This alternative strategy could lead gerrymanderers to draw maps that are deemed unbiased by existing measures of gerrymandering.

For instance, the Efficiency Gap measure quantifies the extent to which a party follows the conventional pack-and-crack strategy by assessing the number of votes 'wasted' by one party relative to the other. If a gerrymanderer adopts a pack-crack-pack strategy instead, the Efficiency Gap may not detect partisan intent. In fact, if there are enough packed districts favoring the gerrymanderer's party, the Efficiency Gap might even suggest that the map benefits the gerrymanderer's opponent. This is because the creation of safe districts is traditionally considered a 'waste' of votes.

References

Abramowitz, Alan I., Brad Alexander, and Matthew Gunning (2006) "Incumbency, Redistricting, and the Decline of Competition in U.S. House Elections," *Journal of Politics*, 68 (1), 75–88.

Baron, David P (1994) "Electoral competition with informed and uninformed voters," *American Political Science Review*, 88 (1), 33–47.

Bouton, Laurent, Micael Castanheira, and Allan Drazen (2018) "A Theory of Small Campaign Contributions,"Technical Report 24413, National Bureau of Economic Research.

Bursztyn, Leonardo, Davide Cantoni, Patricia Funk, Felix Schönenberger, and Noam Yuchtman (2020) "Identifying the Effect of Election Closeness on Voter Turnout: Evidence from Swiss Referenda," Working Paper 23490, National Bureau of Economic Research.

Caughey, Devin, Chris Tausanovitch, and Christopher Warshaw (2017) "Partisan Gerrymandering and the Political Process: Effects on Roll-Call Voting and State Policies," *Election Law Journal: Rules, Politics, and Policy*, 16 (4), 453–469.

Chambers, Christopher P and Alan D Miller (2010) "A measure of bizarreness," *Quarterly Journal of Political Science*, 5 (1), 27–44.

Chen, Jowei and Jonathan Rodden (2013) "Unintentional gerrymandering: Political geography and electoral bias in legislatures," *Quarterly Journal of Political Science*, 8 (3), 239–269.

Cox, Gary W. and Jonathan N. Katz (2007) *Elbridge Gerry's Salamander: The Electoral Consequences of the Reapportionment Revolution*: Cambridge University Press.

DeFord, Daryl R., Nicholas Eubank, and Jonathan Rodden (2022) "Partisan Dislocation: A Precinct-Level Measure of Representation and Gerrymandering," *Political Analysis*, 30, 403–425.

Epstein, Gil S and Shmuel I Nitzan (2006) "Strategic non-binding agreements," *Games and Economic Behavior*, 57 (2), 231–252.

Erikson, Robert S. (1972) "Malapportionment, Gerrymandering, and Party Fortunes in Congressional Elections," *American Political Science Review*, 66 (4), 1234–1245, 10.2307/1957176.

Esteban, Joan and Debraj Ray (2001) "Collective action and the group size paradox," *American Political Science Review*, 95 (3), 663–672.

Friedman, Jeffrey N. and Richard Holden (2009) "The Rising Incumbent Reelection Rate: What's Gerrymandering Got To Do with It?" *Journal of Politics*, 593–611.

Friedman, John N. and Richard Holden (2020) "Optimal Gerrymandering in a competitive environment," *Economic Theory Bulletin*, 8 (2), 347–367.

Friedman, John N. and Richard T. Holden (2008) "Optimal Gerrymandering: Sometimes Pack, but Never Crack," *American Economic Review*, 98 (1), 113–44.

Gelman, Andrew and Gary King (1994) "Enhancing democracy through legislative redistricting," *American Political Science Review*, 88 (3), 541–559.

Genicot, Garance ⓡ Laurent Bouton ⓡ Micael Castanheira (2021) "Electoral Systems and Inequalities in Government Interventions," *Journal of the European Economic Association*.

Gomberg, Andrei, Romans Pancs, and Tridib Sharma (2023) "Electoral Maldistricting," *Working Paper*.

Grofman, Bernard and Gary King (2007) "The Future of Partisan Symmetry as a Judicial Test for Partisan Gerrymandering after LULAC V. Perry," *Election Law Journal*, 6 (1), 2–35.

Gul, Faruk and Wolfgang Pesendorfer (2010) "Strategic Redistricting," *American Economic Review*, 100 (4), 1616–1641.

Herrera, Helios A, Aniol Llorente-Saguer, and Rosemarie Nagel (2014) "The role of information in different types of strategic voting," *Experimental Economics*, 17 (4), 577–609.

Hirshleifer, Jack (1989) "Conflict and rent-seeking success functions: ratio vs difference models of relative success," *Public Choice*, 63 (2), 101–112.

Incerti, Giovanni (2015) "Mixed-Member Electoral Systems and the Objective of Maximizing the Expected Seat Share," *American Journal of Political Science*, 59 (2), 393–408.

Jacobson, Gary C and Samuel Kernell (1985) *Strategy and choice in congressional elections*: Yale University Press.

Jeong, Dahyeon and Ajay Shenoy (2022) "The Targeting and Impact of Partisan Gerrymandering: Evidence from a Legislative Discontinuity," *Review of Economics and Statistics*, 1–47.

Jia, Mofei, Stergios Skaperdas, and Samarth Vaidya (2013) "Contests with small noise and the robustness of the all-pay auction," *Economic Theory*, 52 (3), 833–861.

Jones, Daniel B, Neil Silveus, and Carly Urban (2023) "Partisan Gerrymandering and Turnout," *Journal of Law & Economics*, Forthcoming.

Jones, Daniel B. and Randall Walsh (2018) "How do voters matter? Evidence from US congressional redistricting," *Journal of Public Economics*, 158, 25–47.

Katz, Jonathan N., Gary King, and Elisabeth Rosenblatt (2020) "Theoretical Foundations and Empirical Evaluations of Partisan Fairness in District-Based Democracies," *American Political Science Review*, 114 (1), 164–178.

Kolotilin, Anton and Alexander Wolitzky (2020) "The Economics of Partisan Gerrymandering," UNSW Economics Working Paper No. 2020-12.

Konrad, Kai A (2007) *Strategy and Dynamics in Contests*: Oxford University Press.

Lindbeck, Assar and Jörgen Weibull (1987) "Balanced-budget redistribution as the outcome of political competition," *Public Choice*, 52 (3), 273–297.

Lizzeri, Alessandro and Nicola Persico (2001) "The Provision of Public Goods under Alternative Electoral Incentives," *American Economic Review*, 91 (1), 225–239.

McDonald, Michael D. and Robin E. Best (2015) "Unfair Partisan Gerrymanders in Politics and Law: A Diagnostic Applied to Six Cases," *Election Law Journal*.

Niemi, Richard G, Bernard Grofman, Carl Carlucci, and Thomas Hofeller (1990) "Measuring compactness and the role of a compactness standard in a test for partisan and racial gerrymandering," *The Journal of Politics*, 52 (4), 1155–1181.

Owen, G. and B. Grofman (1988) "Optimal Partisan Gerrymandering," *Political Geography Quarterly*, 7, 5–22.

Persson, Torsten and Guido Tabellini (2000) *Political Economics: Explaining Economic Policy*, MIT Press Books: The MIT Press.

Sherstyuk, Katerina (1998) "How to gerrymander: A formal analysis," *Public Choice*, 27–49.

Shotts, Kenneth W. (2003) "Racial Redistricting's Alleged Perverse Effects: Theory, Data, and "Reality"," *Journal of Politics*, 65 (1), 238–243.

Skaperdas, Stergios and Bernard Grofman (1995) "Modeling negative campaigns," *American Political Science Review*, 89 (1), 49–61.

Snyder, James M (1989) "The effect of party activity on campaign spending," *Legislative Studies Quarterly*, 14 (1), 95–110.

Stashko, Allison (2020) "Crossing the District Line: Border Mismatch and Targeted Redistribution," working paper, University of Utah.

Stephanopoulos, Nicholas O. (2013) "Our Electoral Exceptionalism," *University of Chicago Law Review*, 80 (2).

Stephanopoulos, Nicholas O and Eric M McGhee (2015) "Partisan gerrymandering and the efficiency gap," *U. Chi. L. Rev.*, 82, 831.

Stephanopoulos, Nicholas O. and Christopher Warshaw (2020) "The Impact of Partisan Gerrymandering on Political Parties," *Legislative Studies Quarterly*, 45 (4), 609–643.

Strömberg, David (2004) "Radio's Impact on Public Spending," *The Quarterly Journal of Economics*, 119 (1), 189–221.

_____ (2008) "How the Electoral College Influences Campaigns and Policy: The Probability of Being Florida," *American Economic Review*, 98 (3), 769–807.

Tullock, Gordon (1980) "Efficient Rent Seeking," in *Toward a theory of the rent-seeking society*, 97–112: Texas A&M University Press.

Warrington, Gregory S. (2018) "Quantifying Gerrymandering Using the Vote Distribution," *Election Law Journal: Rules, Politics, and Policy*.

Appendices

A. Vote shares

Given the uniform distribution of η_e, the share of voters from group k that cast their ballot for D is:

$$\sigma_k(\delta) = \tfrac{1}{2} + \phi\left(\bar{\nu}^k - \delta\right).$$

Throughout, we assume that vote shares in each district are "interior":[27]

[A1] Assumption Interior.
We assume that σ_k is strictly between 0 and 1 for any realization of δ by imposing that $\left|\bar{\nu}^P\right| < \frac{1}{2\phi} - \frac{1}{2\gamma}$ for all k, p. Moreover, we assume $\left|\bar{\nu}^P\right| < \frac{1}{2\gamma}$: depending on the realization of δ, σ_k^p can be strictly above or strictly below 1/2.

Next, we derive the vote shares of each party in a given district for any given realization of the aggregate shock δ. With n_{kj} voters of type k in district j, they cast $t_{kj} = n_{kj}\tau_k$ ballots. The variable "t" stands for "turnout". It follows that the number of *votes* in favor of D in district j is the turnout-weighted average of D's support among each group present in the district:

$$\sum_k t_{kj}\sigma_k(\delta) = \frac{\sum_k t_{kj}}{2} + \sum_k t_{kj}\phi \times \left(\bar{\nu}^k - \delta\right),$$

where the last equality derives directly from the value of $\sigma_k(\delta)$. This vote share is increasing in $\bar{\nu}^D$ and $\bar{\nu}^P$, and strictly decreasing in δ.

For the simplicity of exposition, we normalize ϕ to 1 for most of the analysis.

Denoting *total turnout in district* j by $t_j := \sum_k t_{kj}$, following Assumption A1, the probability that D wins district j (at least $1/2$ of the votes) is:

$$\pi_j \quad = \quad \Pr\left(\frac{t_j}{2} + \sum_k t_{kj}\left(\bar{\nu}^k - \delta\right) \geq \frac{t_j}{2}\right) = \Pr\left(\delta \leq \tfrac{\sum_k t_{kj}\bar{\nu}^k}{t_j}\right).$$

Remembering that $\bar{\nu}^k \in \{\bar{\nu}^D, \bar{\nu}^R\}$, that $t_j^D = \sum_d n_{dj}\tau_d$ is the total turnout by Democrat-leaning individuals in district j, and that $t_j^R = \sum_r n_{rj}\tau_r$ is the equivalent total for Republican leaning individuals,

[27] This assumption is but a sufficient condition that simplifies the algebra. It can be relaxed to allow for σ_k to equal 0 or 1. Our results hold as long as $\bar{\nu}^D + \bar{\nu}^R < \frac{1}{2\phi}$ and $\left|\bar{\nu}^P\right| < \frac{1}{2\gamma}$.

the last expression can be expressed as:

$$\pi_j = \Pr\left(\delta \leq \frac{\bar{\nu}^D t_j^D + \bar{\nu}^R t_j^R}{t_j}\right). \tag{A.1}$$

By the distribution of δ and Assumption A1, this simplifies into:

$$\pi_j = \frac{1}{2} + \gamma\left(\frac{t_j^D}{t_j}\bar{\nu}^D + \frac{t_j^R}{t_j}\bar{\nu}^R\right) = \frac{1}{2} + \gamma\left(\frac{t_j^D}{t_j}\Delta + \bar{\nu}^R\right), \text{ where } \Delta \equiv \bar{\nu}^D - \bar{\nu}^R.$$

B. Swaps

Since it only affects vote shares in districts i and j, at the margin, a $k \overset{i}{\rightleftharpoons}_j k'$ swap (defined p12) impacts the expected number of seats by:

$$\left[\frac{\partial \pi_i}{\partial n_{k'i}} - \frac{\partial \pi_i}{\partial n_{ki}}\right] - \left[\frac{\partial \pi_j}{\partial n_{k'j}} - \frac{\partial \pi_j}{\partial n_{kj}}\right]. \tag{B.1}$$

It is easy to check that:

$$\frac{\partial \pi_j}{\partial n_{dj}} = \tau_d \frac{t_j^R}{t_j^2}\,\Delta\gamma \quad \text{and} \quad \frac{\partial \pi_j}{\partial n_{rj}} = -\tau_r \frac{t_j^D}{t_j^2}\,\Delta\gamma.$$

Proof of Property 1

Notice that to establish the concavity or convexity of a $k \overset{i}{\rightleftharpoons}_j k'$ *swap* on $\pi_i + \pi_j$ it suffices to establish whether replacing some type k by type k' in district i has increasing or decreasing return since we are doing the same operation in reverse in district j.

Decreasing returns of DD swaps. The effect of replacing some type d by some type d' in district i is:

$$\Delta\gamma\left(\tau_{d'} - \tau_d\right)\frac{t_i^R}{t_i^2}. \tag{B.2}$$

Let us start with $\tau_{d'} > \tau_{d'}$. As we replace type d by d' in i, t_i increases (while t_i^R remain unchanged). The opposite happens if $\tau_{d'} < \tau_{d'}$. Hence, in both cases, (B.2) decreases.

Increasing returns of RR swaps. The marginal effect on swapping some type r by type r' on π_i is:

$$-\Delta\gamma\left((]\tau_{r'} - \tau_r\right)\frac{t_i^D}{t_i^2}. \tag{B.3}$$

If $\tau_{r'} > (<)\tau_r$, the swap increases (decreases) t_i so the the effect is always increasing.

Concavity or Convexity of RD swaps.

Swapping some r type from district i for some d' type affects π_i as follows:

$$\Delta\gamma \left(\tau_{d'} \frac{t_i^R}{t_i^2} + \tau_r \frac{t_i^D}{t_i^2} \right). \tag{B.4}$$

The second order effect has the same sign as $(\tau_{d'} - \tau_r)\left(\tau_{d'} t_i^R + \tau_r t_i^D\right)$, implying that the marginal effect of the swap is decreasing (increasing) iff $\tau_{d'} > (<)\tau_r$.

Clearly, the same applies for the reverse swap, a **DR swap**. The marginal effect of swapping some d type from district i for some r' type on π_i is:

$$-\Delta\gamma \left(\tau_d \frac{t_i^R}{t_i^2} + \tau_{r'} \frac{t_i^D}{t_i^2} \right). \tag{B.5}$$

The second order effect has the same sign as: $(\tau_{r'} - \tau_d)\left(\tau_d t_i^R + \tau_{r'} t_i^D\right)$, implying that the marginal effect of the swap is decreasing (increasing) iff $\tau_d > (<)\tau_{r'}$.

Profitable Swaps

We are now in a position to derive the marginal effect of each swaps onto the expected number of seats, $\pi_i + \pi_j$. We consider each type of swap in turn:

The effect of $d^{\,i} \rightleftharpoons_j\ r'$ swap on the joint probability of winning, $\pi_i + \pi_j$, is proportional ($\Delta\gamma$ multiplies it) to an expression that can be written in three equivalent manners:

$$\begin{aligned}
\left[\frac{\partial \pi_i}{\partial n_{r'i}} - \frac{\partial \pi_i}{\partial n_{di}}\right] - \left[\frac{\partial \pi_j}{\partial n_{r'j}} - \frac{\partial \pi_j}{\partial n_{dj}}\right] &= (\tau_d - \tau_{r'}) \left[\frac{t_j^R}{t_j^2} - \frac{t_i^R}{t_i^2}\right] + \tau_{r'} \left[\frac{1}{t_j} - \frac{1}{t_i}\right] && \text{(B.6)} \\
&= (\tau_{r'} - \tau_d) \left[\frac{t_j^D}{t_j^2} - \frac{t_i^D}{t_i^2}\right] + \tau_d \left[\frac{1}{t_j} - \frac{1}{t_i}\right] && \text{(B.7)} \\
&= \tau_d \left[\frac{t_j^R}{t_j^2} - \frac{t_i^R}{t_i^2}\right] + \tau_{r'} \left[\frac{t_j^D}{t_j^2} - \frac{t_i^D}{t_i^2}\right] && \text{(B.8)}
\end{aligned}$$

The effect of $r \overset{i}{\rightleftharpoons}_j d'$ swap on the joint probability of winning, $\pi_i + \pi_j$, is similarly proportional to any of the following three equivalent expressions:

$$(\tau_{d'} - \tau_r)\left[\frac{t_i^R}{t_i^2}\frac{t_j^R}{t_j^2}\right] + \tau_r\left[\frac{1}{t_i} - \frac{1}{t_j}\right] \tag{B.9}$$

$$(\tau_r - \tau_{d'})\left[\frac{t_i^D}{t_i^2} - \frac{t_j^D}{t_j^2}\right] + \tau_{d'}\left[\frac{1}{t_i} - \frac{1}{t_j}\right] \tag{B.10}$$

$$\tau_{d'}\left[\frac{t_i^R}{t_i^2} - \frac{t_j^R}{t_j^2}\right] + \tau_r\left[\frac{t_i^D}{t_i^2} - \frac{t_j^D}{t_j^2}\right] \tag{B.11}$$

The effect of $d \overset{i}{\rightleftharpoons}_j d'$ swap on the joint probability of winning, $\pi_i + \pi_j$ is proportional to:

$$(\tau_{d'} - \tau_d)\left[\frac{t_i^R}{t_i^2} - \frac{t_j^R}{t_j^2}\right] \tag{B.12}$$

The effect of $r \overset{i}{\rightleftharpoons}_j r'$ swap on the joint probability of winning, $\pi_i + \pi_j$ is proportional to:

$$(\tau_{r'} - \tau_r)\left[\frac{t_j^D}{t_j^2} - \frac{t_i^D}{t_i^2}\right] \tag{B.13}$$

C. Proofs: Main Results

Proof of Proposition 1.
Without loss of generality, we assume that the gerrymanderer is a Democrat.

Proof of part(1). Consider district i with $n_{di}, n_{ri} > 0$ and district j with $n_{d'j}, n_{r'j} > 0$. That is, there are some democrats and some republicans in both districts, with index k associated with those initially in district i and index k' associated with those initially in district j. Let $\max\{\tau_d, \tau_{d'}\} < \min\{\tau_r, \tau_{r'}\}$: partisans have strictly lower turnout rates than opponents.

Without loss of generality, let types $\tau_{d'}$ in district j have the lowest turnout: $\tau_{d'} \leq \tau_d$. This means there are up to six cases to consider: (I-III) $\tau_{d'} = \tau_d < \tau_r \gtreqless \tau_{r'}$, and (IV-VI) $\tau_{d'} < \tau_d < \tau_r \gtreqless \tau_{r'}$.

In an optimal districting, none of the four possible swaps ($d \overset{i}{\rightleftharpoons}_j d'$, $r \overset{i}{\rightleftharpoons}_j r'$, $d \overset{i}{\rightleftharpoons}_j r'$, and $r \overset{i}{\rightleftharpoons}_j d'$) may have a strictly positive marginal effect on $\pi_i + \pi_j$. Starting with $d \overset{i}{\rightleftharpoons}_j d'$, (B.12) must be non-positive:

$$(\tau_d - \tau_{d'}) \left[\frac{t_j^R}{t_j^2} - \frac{t_i^R}{t_i^2} \right] \leq 0. \tag{NoDD}$$

Next, consider $r \overset{i}{\rightleftharpoons}_j r'$ swaps: for $\tau_r \neq \tau_{r'}$, (B.13) must be strictly negative:

$$(\tau_r - \tau_{r'}) \left[\frac{t_i^D}{t_i^2} - \frac{t_j^D}{t_j^2} \right] < 0. \tag{NoRR}$$

The strict inequality follows from the convexity of RR swaps in Property 1.

Finally, the two possible DR swaps must have a negative impact on $\pi_i + \pi_j$. The effect of a $d \overset{i}{\rightleftharpoons}_j r'$ swap must be negative:

$$(\tau_{r'} - \tau_d) \left[\frac{t_j^D}{t_j^2} - \frac{t_i^D}{t_i^2} \right] + \tau_d \left[\frac{1}{t_j} - \frac{1}{t_i} \right] \leq 0, \tag{NoDR}$$

and the effect of a $r \overset{i}{\rightleftharpoons}_j d'$ swap must be negative:

$$(\tau_r - \tau_{d'}) \left[\frac{t_i^D}{t_i^2} - \frac{t_j^D}{t_j^2} \right] + \tau_{d'} \left[\frac{1}{t_i} - \frac{1}{t_j} \right] \leq 0. \tag{NoRD}$$

Case I: $\tau_d = \tau_{d'} < \tau_r = \tau_{r'}$. It is easy to show that this case cannot be an optimal districting as the gerrymanderer can strictly increase their payoff through some $d \overset{i}{\rightleftharpoons}_j r'$ or $r \overset{i}{\rightleftharpoons}_j d'$ swaps. To see this recall that, by Property 1(3), the effect of either swap is convex since republicans have the turnout rate advantage. This implies that: (i) if both (NoDR) and (NoRD) are satisfied with equality, the expected number of seats going to party D must be at a local minimum, and (ii) if either (NoDR) or (NoRD) is non zero, then one must be strictly positive since (NoDR) and (NoRD) are the opposite of each other.

Case II. $\tau_{d'} = \tau_d < \tau_{r'} < \tau_r$. We note that (NoRR) only holds if: $\frac{t_i^D}{t_i^2} - \frac{t_j^D}{t_j^2} < 0$. This also implies that *all* R types in district i must have a higher turnout rate than all R types in district j, otherwise there is at least one profitable $r \overset{i}{\rightleftharpoons}_j r'$ swap. Note also that all D types must have the same turnout rate in this case, otherwise Case V or VI applies. The value of a $d \overset{i}{\rightleftharpoons}_j r'$ in (NoDR) can therefore be decomposed into a first term that is strictly positive, and a second term that must be sufficiently negative for the $d \overset{i}{\rightleftharpoons}_j r'$ to be non-profitable. This requires that $t_i < t_j$. Given the aforementioned conditions on turnout rates, $t_i < t_j$ requires that there are more Democrat voters in i than in j, *i.e.* $t_i^D/t_i > t_j^D/t_j$. Put

together, these two conditions in turn require that $\frac{t_i^D}{t_i^2} - \frac{t_j^D}{t_j^2} > 0$, meaning that (NoRR) must be positive and hence that Case II cannot be optimal.

Case III. $\tau_{d'} = \tau_d < \tau_r < \tau_{r'}$. This is the same as Case II except for the fact that $\tau_r < \tau_{r'}$. Relabeling the districts i and j proves that Case III cannot be an optimal redistricting.

Case IV. $\tau_{d'} < \tau_d < \tau_{r'} = \tau_r$. If inequality (NoDD) is strict, the fact that (NoDD) must hold for all D types implies that $\tau_d \geq \tau_{d'}$ for all D types in districts i and j. Note also that all R types must have the same turnout rate in districts i and j, otherwise case V or VI applies. It must then be that $t_i > t_j$. For this not to be the case, district j would need to have strictly fewer D types than district i (since all D types in i have higher turnout rates than all D types in j and all R types have the same turnout rate). This implies that $t_j^D < t_i^D$, which, together with $t_j \geq t_i$, implies $t_j^D/t_j^2 < t_i^D/t_i^2$. However, $t_j^D/t_j^2 < t_i^D/t_i^2$ and $t_j \geq t_i$ violate condition (NoRD). Thus, it must be that $t_i > t_j$. Given $t_i > t_j$ and that (NoDD) is strict, along with (B.6), we have that a $d^i \rightleftharpoons_j r'$ swap must be positive (condition (NoDR) cannot be satisfied).

If inequality (NoDD) is not strict, then $t_j^R/t_j^2 = t_i^R/t_i^2$. Note that $t_j^R/t_j^2 = t_i^R/t_i^2$ implies $t_j^D/t_j^2 - t_i^D/t_i^2 = 1/t_j - 1/t_i$. Thus, inequalities (NoDR) and (NoRD) simplify to:

$$\left[\frac{t_j^D}{t_j^2} - \frac{t_i^D}{t_i^2}\right] \leq 0 \quad \text{and} \quad \left[\frac{t_i^D}{t_i^2} - \frac{t_j^D}{t_j^2}\right] \leq 0. \tag{C.1}$$

Given that these inequalities are the opposite of each other, they can only be jointly satisfied if equal to zero. But, for the ordering of turnout rates under consideration, we have from Property 1(3) that the effect of either swap is convex. Hence, we would be at a local minimum.

It follows that Case IV can not be an optimal districting.

Case V. $\tau_{d'} < \tau_d < \tau_{r'} < \tau_r$. With $\tau_r > \tau_{r'}$, (NoRR) requires

$$\frac{t_i^D}{t_i^2} - \frac{t_j^D}{t_j^2} < 0. \tag{C.2}$$

whereas with $\tau_d > \tau_{d'}$, (NoDD) requires

$$\frac{t_i^R}{t_i^2} - \frac{t_j^R}{t_j^2} \geq 0. \tag{C.3}$$

Given (C.2) and that (NoRR) must hold for all R types, then all R types present in i must have a higher turnout rate than any of the R types present in district j. If inequality (C.3) is strict, then by (NoDD) all

D types present in i must have a higher turnout rate than any of the D types present in district j. This in turn implies that $t_i > t_j$. For this not to be the case, one would need to have more D types in district i than in district j. Since D types in district i have a higher turnout rate than D types in district j, this implies $t_i^D > t_j^D$. Along with the assumption that $t_j > t_i$, this implies $t_i^D/t_i^2 > t_j^D/t_j^2$, which means the effect of an $r \overset{i}{\rightleftharpoons}_j r'$ is positive, or condition (NoRR) is violated. Thus, we must have $t_i > t_j$. In that case, (NoDR) must be positive.

If (C.3) holds with equality, then $\frac{t_j^R}{t_j^2} - \frac{t_i^R}{t_i^2} = 0$. As in Case IV, this implies that expressions (NoDR) and (NoRD) cannot both be true. Thus, Case V can not be an optimal districting.

Case VI. $\tau_{d'} < \tau_d < \tau_r < \tau_{r'}$. This is the same as Case V except for the fact that $\tau_r < \tau_{r'}$. Hence, condition (NoRR) now requires:

$$\frac{t_i^D}{t_i^2} - \frac{t_j^D}{t_j^2} > 0. \tag{C.4}$$

Given $\tau_{d'} < \tau_d$, (NoDD) still requires:

$$\frac{t_i^R}{t_i^2} - \frac{t_j^R}{t_j^2} \geq 0. \tag{C.5}$$

Adding up inequalities (C.4) and (C.5) implies that $t_i < t_j$. Together, the latter inequality and (C.4) imply that (NoRD) is not satisfied.

Together, cases I-VI show that there is no configuration of turnout rates for which it may be optimal to maintain a pair of districts in which low-turnout partisans (D types) and high-turnout opponents (R types) co-exist.

Proof of part (2). Take any two districts i and j such that $\min\{\tau_d, \tau_{d'}\} > \max\{\tau_r, \tau_{r'}\}$ for any d, r types in district i and d', r' types in district j. We cannot have $\sum_{r\in\mathcal{R}} n_{ri} = 0$ and $\sum_{d\in\mathcal{D}} n_{dj} = 0$.

Let us proceed by contradiction and assume that both $\sum_{r\in\mathcal{R}} n_{ri} = 0$ and $\sum_{d\in\mathcal{D}} n_{dj} = 0$ (that is, we start from a situation with a packed D district and a packed R district). Since R-types have by assumption lower-turnout rates than D-types, this implies that $t_i > t_j$. Since $\sum_{r\in\mathcal{R}} n_{ri} = 0$, then we also have that $t_i^R/t_i^2 = 0$. Hence, (B.6) is strictly positive: it is strictly profitable to perform a $d \overset{i}{\rightleftharpoons}_j r'$ swap, a contradiction. ■

Proof of Proposition 2. Assume that the claim is incorrect. There is a district i with two types of republicans $r, \widetilde{r}$ and a district j with two types of republicans r' (potentially equal to r) and $\widetilde{r}'$ (potentially equal to $\widetilde{r}$), such that $\tau_r > \tau_{\widetilde{r}'}$ and $\tau_{r'} > \tau_{\widetilde{r}}$.

At the optimal districting, a swap of r types from district i and $\widetilde{r}'$ types from district j cannot increase the gerrymander's objective. Using (B.13) this implies that $\frac{t_i^D}{t_i^2} \leq \frac{t_j^D}{t_j^2}$ Similarly, a swap of r' types from district j and $\widetilde{r}$ types from district i cannot increase the gerrymander's objective implies that $\frac{t_i^D}{t_i^2} \geq \frac{t_j^D}{t_j^2}$.

Together it means that $\frac{t_i^D}{t_i^2} = \frac{t_j^D}{t_j^2}$ and both types of swaps have a zero marginal effect on the expected number of seats won.

Recall from Property 1 that the gerrymander's objective is convex in the swaps of Republicans of different turnout rates across districts. It follows that a discrete swap in one direction or the other would increase the expected number of seats won. ■

Proof of Proposition 3. Suppose not. Then there exists a district i with types d and r and a district j with types d' and r' where $\tau_{r'} < \tau_r$ and $\tau_d < \tau_{d'}$. In an optimal districting, the effect of a $d \overset{i}{\rightleftharpoons}_j d'$ swap (equation B.12) must be non-positive. Given that $\tau_{d'} > \tau_d$, this requires:

$$\frac{t_i^R}{t_i^2} - \frac{t_j^R}{t_j^2} \leq 0. \tag{C.6}$$

The effect of a $r \overset{i}{\rightleftharpoons}_j r'$ swap (equation B.13) in an optimal districting must be strictly negative. Given $\tau_r > \tau_{r'}$, this requires

$$\frac{t_i^D}{t_i^2} - \frac{t_j^D}{t_j^2} < 0. \tag{C.7}$$

From (B.8), we see that these two conditions imply that there would be a strictly profitable $d \overset{i}{\rightleftharpoons}_j r'$ swap contradicting the optimality of the considered map. ■

Proof of Proposition 4.

Without loss of generality, the gerrymanderer is D. We start by showing that if r types are packed, then all types r' such that $\tau_{r'} > \tau_r$ must also be packed. Suppose that there exists a packed district i with $n_{ri} > 0$. District i contains no d types, so $\frac{t_i^D}{t_i^2} = 0$. In district j, there is some r' with $n_{r'j} > 0$ and $\tau_{r'} > \tau_r$. From equation (B.13), the effect of a $r \overset{i}{\rightleftharpoons}_j r'$ swap between packed district i and cracked j has the same sign as:

$$(\tau_r - \tau_{r'}) \left[\frac{t_i^D}{t_i^2} - \frac{t_j^D}{t_j^2} \right] = (\tau_{r'} - \tau_r) \frac{t_j^D}{t_j^2} \geq 0.$$

Hence, the swap is strictly profitable as soon as $t_j^D > 0$: the allocation may only be an equilibrium if $t_j^D = 0$, that is if j is also a packed district.

Following the same logic, if d types are packed, then all types d' such that $\tau_{d'} < \tau_d$ must also be packed. ■

Proof of Proposition 5. The claim is that all cracked districts can be ordered by increasing turnout rates and decreasing winning probabilities.

By Proposition 2, all mixed districts can be ordered by non decreasing Republican turnout rate. That is, we can order districts so that for any two districts i and j with $i < j$, we must have $\max_{r \in i} \tau_r \leq \min_{r' \in j} \tau_{r'}$, where "$r \in i$" is short for "there are at least some types r in district i: $n_{ri} > 0$".

Moreover, Proposition 3 tells us that we have assortative matching among cracked districts. It follows that we can order districts so that for any two districts i and j with $i < j$: $\max_{d \in i} \tau_d \leq \min_{d' \in j} \tau_{d'}$.

It remains to prove the following:

CLAIM: $\pi_i(\mathbf{n}_i) \geq \pi_j(\mathbf{n}_j)$ or, in other words, $\frac{t_i^D}{t_i} \geq \frac{t_j^D}{t_j}$.

Assume not and instead that $\pi_i(\mathbf{n}_i) < \pi_j(\mathbf{n}_j)$ or, equivalently, that $\frac{t_i^D}{t_i} < \frac{t_j^D}{t_j}$.

CASE 1: there exist two types r' in j and r in i with different turnout rates, *i.e.* $\tau_{r'} > \tau_r$. From (NoRR), we need $\frac{t_i^D}{t_i^2} \geq \frac{t_j^D}{t_j^2}$ so that an $r \overset{i}{\rightleftharpoons}_j r'$ swap is not profitable. Since $\frac{t_i^D}{t_i} < \frac{t_j^D}{t_j}$ then it must be that $\frac{1}{t_i} > \frac{1}{t_j}$. Given that $\frac{t_i^D}{t_i^2} \geq \frac{t_j^D}{t_j^2}$, from (NoRD) (or equivalently (B.11)) we need $\frac{t_j^R}{t_j^2} \geq \frac{t_i^R}{t_i^2}$ so that an $r \overset{i}{\rightleftharpoons}_j d'$ swap is not profitable. Since $\frac{1}{t_i} > \frac{1}{t_j}$, it must be that $\frac{t_j^R}{t_j} > \frac{t_i^R}{t_i}$, and therefore that $\frac{t_j^D}{t_j} < \frac{t_i^D}{t_i}$, a contradiction. We must thus have that $\frac{t_i^D}{t_i} \geq \frac{t_j^D}{t_j}$.

CASE 2: If there is only one type of republicans $r = r'$ in both districts but (at least) two types of democrats d' in district j and d in district i so that $\tau_d < \tau_{d'}$. From (NoDD) (or equivalently (B.12)), we need $\frac{t_i^R}{t_i^2} \leq \frac{t_j^R}{t_j^2}$ so that a $d \overset{i}{\rightleftharpoons}_j d'$ swap is not profitable. Then (B.8) requires that $\frac{t_j^D}{t_j^2} \leq \frac{t_i^D}{t_i^2}$, otherwise a $d \overset{i}{\rightleftharpoons}_j r'$ swap would be profitable. Since $\frac{t_i^D}{t_i} < \frac{t_j^D}{t_j}$, it must then be that $\frac{1}{t_i} > \frac{1}{t_j}$. But then $\frac{t_i^R}{t_i^2} \leq \frac{t_j^R}{t_j^2}$ implies that $\frac{t_i^R}{t_i} < \frac{t_j^R}{t_j}$ (*i.e.* $\frac{t_i^D}{t_i} > \frac{t_j^D}{t_j}$) a contradiction.

CASE 3: If there is only one type of republicans $r = r'$ and one type of democrats $d = d'$ in both districts, we can order them such that j has the (weakly) lower probability of winning. ■

C.1. Robust maps

To further describe the optimal map, we want to isolate the effects of the population structure, and abstract from the constraints imposed by a possibly small number of districts. Related models rule out such constraints by assuming a continuum of districts (Kolotilin and Wolitzky (2020), Friedman and Holden 2008, 2020, Gul and Pesendorfer 2010). Instead, we introduce the concept of a *robust map*, which is insensitive to the number of districts.

To understand the issue, consider a case in which there are only low-turnout supporters and high-turnout opponents. By Proposition 1, the gerrymanderer avoids mixing them in a same district. Now, imagine there are 75% supporters and 25% opponents but there are only two districts. Then, the gerrymanderer can only create one packed district filled with supporters. In the second district, they are forced to mix the remaining 25% of supporters with the 25% of opponents.

Now, clone this map by doubling the number of districts, while keeping the population unchanged. The gerrymanderer can now create four packed districts, three filled with supporters, and only one with opponents. In this example, four districts are sufficient to describe an optimal map. Moving beyond this simple example, we allow for the indefinite *cloning* of a map:

Definition: *The cloning of a map consists in the replication of each district j, creating $c \in \mathbb{N}_0$ additional districts with the same population structure as the original one.*

It is straightforward to check that cloning may only serve the gerrymanderer: the cloned map results, by definition, in the same payoff as the original map. Yet, it may also open the door to additional swaps that were initially not available. Instead, when all room for optimization has been exhausted, further cloning becomes immaterial. For instance, increasing the number of districts from 4 to 8 in the example does not affect payoffs. This defines robust maps:

Definition: *A map is robust if, for any $c \in \mathbb{N}_0$, cloning does not allow a gerrymanderer to further increase their expected seat share.*

Robustness allows us to determine sufficient conditions for the existence –or co-existence– of the packed-D districts, cracked districts, and packed-R districts, *i.e.* a pack-crack-pack pattern, as identified in Proposition 4. In particular, (1) whenever the highest turnout rate of democrats is higher than the lowest turnout rate of republicans, a robust map must display at least one cracked district; (2) opponents who turn out more than any supporters will necessarily be allocated to packed-R districts; (3) supporters who turn out less than any opponent will be allocated to packed-D districts.

This leads to Proposition 6:

Proof of Proposition 6.

Part (1). Assume not and that instead all districts are packed, some with supporters and some with opponents. Note that the payoff of any two fully packed maps is equal. We can thus clone our fully packed map c times so that $c \times \min\{n_{\bar{d}}, n_{\underline{r}}\} \geq 1/J$, where $1/J$ is the mass of voters needed to populate a district without changing the gerrymanderer's payoff. Moreover, the gerrymanderer's payoff under this fully packed map is the same as the payoff associated with another fully packed map in which at least one district is only composed of partisans with turnout rates $\tau_{\bar{d}}$ and at least one district is only composed of opponents with turnout $\tau_{\underline{r}} (< \tau_{\bar{d}})$. However, this leads to a contradiction: Proposition 1(2) tells us that the expected seat share can be strictly increased by mixing these high-turnout supporters with these low-turnout opponents. This establishes that:

$$\tau_{d^*} < \tau_{\bar{d}} \text{ and } \tau_{r^*} > \tau_{\underline{r}},$$

i.e. there must be a strictly positive number of mixed districts.

Part (2). To prove point (2), we show that if $\tau_{\bar{d}} < \tau_{\bar{r}}$ then $\tau_{r^*} < \tau_{\bar{r}}$, *i.e.* any robust map has at least one district packed with opponents: In the absence of a packed district, there must be some mixed district j with $n_{\bar{r}j} > 0$. Cloning the map ensures that there are at least two cracked districts that mix $\tau_{\bar{r}}$ with some lower turnout D types. By Proposition 1(1), this cannot happen in an optimal map.

Part (3). We show that if $\tau_{\underline{d}} < \tau_{\underline{r}}$ then $\tau_{d^*} > \tau_{\underline{d}}$. In the absence of a packed district only populated with $\tau_d < \tau_{\underline{r}}$, there must be some cracked district i that mixes low-turnout supporters with higher-turnout opponents: $\exists i$ with $n_{\underline{d}i} > 0$ and $n_{ri} > 0$, s.t. $\tau_r > \tau_{\underline{d}}$. Cloning the map, there are at least two such districts. By Proposition 1(1), this cannot happen in an optimal map. ■

D. Extension: Endogenous Turnout

We extend the model to a setup in which turnout rate is the product of (a) the share of eligible voters τ_k in group k and (b) the turnout decision σ_k of those who are eligible: $\tau_k \times \sigma_k$.

Those eligible to vote turn out only when their expected benefit of voting outweighs the cost, in line with the calculus of voting tradition (as described in, for example, Riker and Ordeshook 1968). The expected benefit depends on both the utility differential between the parties and a function $p_k(\mathbf{n_j})$, which maps the composition of district j to a benefit of voting for eligible voters in group k. The benefit of voting

can vary, for example, with the partisan composition in district j. Moreover, population groups can differ in how sensitive they are to district-level characteristics. Voting also entails a cost c_k, which can be group-specific. We thus have that an eligible individual i from group k in district j votes for D if

$$(\bar{\nu}_k - \eta_{ek} - \delta)\, p_k(\mathbf{n_j}) \geq c_k.$$

They vote for R if:

$$c_k \leq -(\bar{\nu}_k - \eta_{ek} - \delta)\, p_k(\mathbf{n_j}),$$

and they abstain otherwise.

For a given value of δ, the share of eligible individuals who turn out to vote for D is then:

$$\sigma_{kj}^{D}(\delta) = \frac{1}{2} + \phi\left(\bar{\nu}_k - \delta - \frac{c_k}{p_k(\mathbf{n_j})}\right)$$

whereas the share who votes for R is:

$$\sigma_{kj}^{R}(\delta) = \frac{1}{2} - \phi\left(\bar{\nu}_k - \delta + \frac{c_k}{p_k(\mathbf{n_j})}\right).$$

It follows that the actual number of votes for D, given δ is:

$$t_j^D = \sum_k n_{kj} \cdot \tau_k \cdot \sigma_{kj}^{D}(\delta) = \frac{\sum_k n_{kj} \cdot \tau_k}{2} + \phi \sum_k n_{kj} \cdot \tau_k \cdot \left(\bar{\nu}_k - \delta - \frac{c_k}{p_k(\mathbf{n_j})}\right)$$

and similarly for party R:

$$t_j^R = \frac{\sum_k n_{kj} \cdot \tau_k}{2} - \phi \sum_k n_{kj} \cdot \tau_k \cdot \left(\bar{\nu}_k - \delta + \frac{c_k}{p_k(\mathbf{n_j})}\right).$$

Hence, the probability that D wins district j is:

$$\pi_j = \Pr\left(t_j^D \geq t_j^R\right) = \Pr\left(\delta \leq \frac{t_j^D \bar{\nu}_D + t_j^R \bar{\nu}_R}{t_j}\right),$$

which is the same as in the baseline model, where turnout rates are exogenous.

E. Extension: Single-Peaked Aggregate Shock

The objective of the gerrymanderer is to maximize the expected number of districts won, $\sum_j \pi_j$, where the expected probability of winning district j is $\pi_j = \Gamma\left(\hat{\delta}_j\right)$, where $\hat{\delta}_j = \frac{t_j^D \bar{\nu}^D + t_j^R \bar{\nu}^R}{t_j}$ is the pivotal aggregate shock such that the party wins the district.

Traditional Pack-and-Crack

Assume that all turnout rates are the same. In this case, for each district j, $\hat{\delta}_j$ is the average partisanship of the district. If Γ was linear (uniform distribution of δ), the gerrymanderer would then maximize a sum of average partisanship and therefore be indifferent between all the maps.

In the rest of this appendix, we instead focus on the case of a non-uniform distribution for δ. Assume instead that $\frac{\partial \Gamma(\delta)}{\partial \delta} = \gamma(\delta)$ is single-peaked around $\delta = 0$.

If $\bar{\nu}^R < 0 < \bar{\nu}^D$, π_j is S-shaped in the average partisanship of the district, and as shown by Owen and Grofman (1988) and Kolotilin and Wolitzky (2020) we get the traditional pack-and-crack strategy. There are at most two types of districts, some packed with opponents and some cracked, all with the same mix of supporters and opponents.

Extended Model

Let us keep our single-peaked assumption ($\bar{\nu}_D > 0 > \bar{\nu}_R$) and reintroduce the heterogeneity in turnout rates across groups. This allows us to explore the robustness of our results to a non-uniform distribution of the aggregate shock δ.

Properties of the Swaps

We first consider how the single-peakedness of the distribution of aggregate shocks affects the concavity or convexity of the swaps (i.e., we revisit Property 1).

RD swaps. The effect of simultaneously increasing $n^i_{d'}$ and decreasing n^i_r on $\pi_i = \Gamma\left(\hat{\delta}_i\right)$ is given by:

$$\gamma_i \Delta \frac{\tau_{d'} t_j^R + \tau_r t_i^D}{t_i^2}.$$

where $\gamma_i = \gamma\left(\hat{\delta}_i\right)$. The second order effect has the same sign as:

$$\gamma_i' \Delta \frac{\left(\tau_{d'} t_i^R + \tau_r t_i^D\right)}{t_i} - 2\gamma_i \cdot (\tau_{d'} - \tau_r)$$

where $\gamma_i' = \gamma'\left(\hat{\delta}_i\right)$.

The effects described in Property 1 are captured by the second term: the effects of the swap are concave (convex) if $\tau_d > (<)\tau_r$. The additional effect brought about by a non-uniform distribution of δ is

captured by the first term, through γ_i'. If $\hat{\delta}_i > 0$ then $\gamma_i' < 0$, which means that this added effect is concave. If $\hat{\delta}_i < 0$ then $\gamma' > 0$ and the added effect is convex.

Following similar steps, we can show that (i) the concavity of the DD swaps from Property 1 are reinforced when $\gamma' < 0$ ($\hat{\delta}_i > 0$), but not when $\gamma' > 0$ ($\hat{\delta}_i < 0$), and (ii) The convexity of the RR swaps from Property 1 are reinforced when $\gamma' > 0$ ($\hat{\delta}_i < 0$), but not when $\gamma' < 0$ ($\hat{\delta}_i > 0$).

Implications

In this extended model, both the traditional pack and crack incentives and our turnout differential incentives are present.

We have seen that the concavity of Γ in strong districts ($\hat{\delta}_i > 0$) reinforces the concavity of the DD swaps and of the DR swaps when $\tau_d > \tau_{r'}$. In contrast, the convexity of Γ in weak districts ($\hat{\delta}_i < 0$) goes against it.

Broadly speaking, the main effect of a single-peaked distribution for the aggregate shock is that it reduces the incentives for the gerrymanderer to create multiple cracked districts that they expect to lose in expectation. Instead, the gerrymanderer may prefer to abandon one of those districts (and transform it into a packed R district) in order to free up some supporters to reinforce the other districts. This follows from the convexity of Γ in weak districts. This force may also lead the gerrymanderer to add lower turnout supporters to a cracked district to reinforce it. This pattern is illustrated in the example below.

Nonetheless, we can show that our main results hold, at least qualitatively, in this extended setup. First, the pack-crack-pack structure of the optimal map holds. If the gerrymanderer creates any packed D district, it must be composed of supporters with the lowest turnout rate(s). Otherwise, it is easy to show that there necessarily exists a profitable DD swap. Similarly, any packed R district must be composed of opponents with the highest turnout rate(s).

Second, we can also prove that there are cases in which the three types of districts in the pack-crack-pack structure coexist in an optimal map. Our turnout differential incentives and the traditional pack-and-crack incentives generate two different reasons to pack high turnout opponents, but these incentives fight each other when it comes to packing low turnout supporters. Which effect dominates then depend on the curvature of the distribution compared to the difference in turnout rates. To illustrate this consider two districts: a district i that is packed D with one type of democrat d and a district j that is mixed. The

profitability of a $d \overset{i}{\rightleftharpoons}_j r'$ then depends on the sign of:

$$\gamma_j \frac{\tau_d t_j^R + \tau_{r'} t_j^D}{t_j^2} - \gamma_i \frac{\tau_{r'}}{t_i^D}.$$

It is easy to verify that a sufficiently low τ_d is a sufficient condition for the swap not to be profitable (i.e., the gerrymanderer does not want to unpack district i).

Finally, as in the general model, we can order districts by declining probability of winning and can show that if district j has a lower probability of winning than district j both democrats and republicans have higher turnout rates than their counterparts in district i.[28]

An Example

Here is an example to illustrate that, in some cases, all types of districts in the pack-crack-pack structure can coexist in an optimal map.

Let $\bar{\nu}_D = 1 = -\bar{\nu}_R$. Assume that the population is partitioned in four groups: $3/8$ low turnout republicans $\underline{r}$ ($\tau_{\underline{r}} = 0.4$), $1/4$ high turnout republicans $\bar{r}$ ($\tau_{\bar{r}} = 0.8$), $1/4$ of low turnout democrats $\underline{d}$ ($\tau_{\underline{d}} = 0.2$), and $1/8$ high turnout democrats $\bar{d}$ ($\tau_{\bar{d}} = 0.5$).

Figure E.1 contrasts the optimal maps for the case of a uniform distribution of the aggregate shock ($U[-5, 5]$) and the case of a normal distribution ($N(0, 1)$). In both cases, low turnout democrats and high turnout republicans are packed. Under the uniform, two cracked districts are created mixing the high turnout democrats and the low turnout republicans despite the fact that these districts have a low expected probability of winning. Under the normal instead, the gerrymanderer prefers creating only one such mixed district to strengthen its lead in that district.

[28] Proof available upon request.

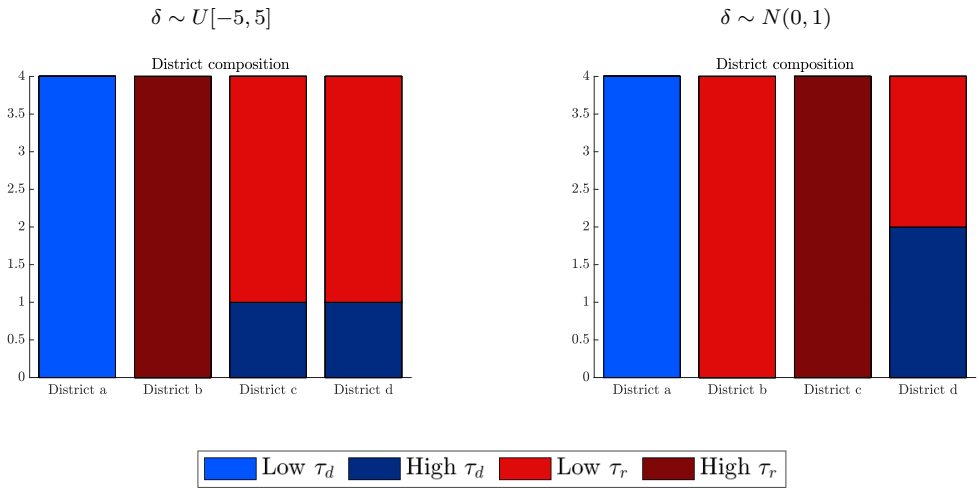


FIGURE E.1. Pack-Crack-Pack.

F. Data Appendix

F.1. Dataset

We compile precinct-level political and demographic data from several different sources. We then use geospatial data to assign the precinct-level characteristics to congressional districts under Democrat and Republican proposals from the 2020 redistricting cycle.[29]

The ability to compare how a precinct is treated under a Democrat versus Republican proposal is key for our analysis. Therefore, we limit attention to states in which we observe a proposed map from both political parties. Detailed data for redistricting proposals, and not just the final map, became widely available only after the 2020 redistricting cycle. The maps themselves were drawn by redistricting committees, party caucuses, or individual representatives (see Appendix F.2).

[29]We focus on U.S. congressional districts rather than state legislative districts to avoid incumbency gerrymandering, as much as possible. Since most maps are drawn by state legislators, incentives for incumbency gerrymandering are relatively strong at the state level. Given that most states have only a small number of congressional districts, an advantage of our theoretical framework is that we consider a finite number of districts.

We further restrict attention to states with high quality information on party affiliation among registered voters. Party affiliation is a key input that we use to determine the partisan lean of a precinct or legislative district.[30] In total, we collect data for 20 redistricting proposals from 10 states (Florida, Kansas, Louisiana, Maryland, Nebraska, Nevada, New Mexico, New York, North Carolina, and Pennsylvania), covering 113 congressional districts.

At the precinct-level, we compile information on demographics, voter registration, and presidential election returns from 2016 and 2020. Precincts are geographic areas used to administer elections. They are the smallest possible geographic unit to observe election returns. The precinct-level data come from multiple sources. We use 2020 Census data to measure the total population. We use American Community Survey data to measure age, race/ethnicity, sex, income, and education. The Redistricting Data Hub compiles voter registration data, election returns, and Citizen Voting Age Population (CVAP, the number of citizens aged 18 and older).[31]

F.2. Redistricting Proposals

We collect redistricting proposals from the 2020 redistricting cycle for all states in which we observe both a Democrat and a Republican proposal. The sources of the proposals vary from state to state. In most states, we observe a proposal from the redistricting committee of the state legislature. Committees charged with redistricting are typically aligned with the majority party in the legislature. We therefore assign a legislative committee map to the majority party. For the minority party, we use proposals from party caucuses and from individual representatives. If there are multiple proposals for the minority party, we use the proposal from a caucus over an individual. New York is an exception because an independent committee issued two partisan maps for congressional districts. See Table F.1 for details. We compare the map proposed by the Republican-nominated members of the independent redistricting committee to the map proposed by the Democrats of the New York state legislature. Geospatial data for each proposal were gathered from state legislature websites or by request to state legislatures. The Geospatial data allows us to map each proposal and measure how the proposed districts intersect with existing precincts.

[30] Party affiliation is a useful proxy because it is relatively stable over time, unlike election returns. Using vote shares would be problematic because vote shares are also a function of turnout, and possibly endogenous to redistricting. We use only states where voters are required to register with a party in order to participate in the party's primary elections. We expect the quality of party affiliation data to be higher in these states than in states where party registration is optional or imputed.

[31] The Redistricting Data hub geocodes and validates precinct-level data. Their source for voter registration data is L2, their source for election returns is the Voting and Election Science Team (VEST), and their source for CVAP is the American Community Survey.

TABLE F.1. Redistricting Proposals

State	Party	Author	Proposal Name	Notes
FL	Dem	State Senator Darryl Rouson	S019C8062	Of the proposals from Democrats, this was submitted most recently.
FL	Rep	Governor Ron DeSantis	PP000C0109109	Governor DeSantis rejected the legislature's plan and proposed this one, which was later enacted and challenged in state court.
KS	Dem	State Senator Dinah Sykes	United	Senator Sykes is the minority party leader.
KS	Rep	Kansas House Committee on Redistricting	Ad Astra 2	The democratic governor vetoed this proposal and the legislature overrode the veto. The Kansas Supreme Court ultimately rejected the map.
LA	Dem	State Senator Gary Smith	SB11	Several proposals were submitted by Democrats in a short time span. This plan was from the most senior member of the state legislature.
LA	Rep	State Representative Clay Schexnayder	House Bill 1	The Democratic governor vetoed this proposal and the legislature overrode the veto. An appeals court rejected the map, but the U.S. Supreme Court allowed it to be used in 2022 elections.
MD	Dem	Legislative Redistricting Advisory Commission	LRAC Final Recommended Congressional Plan	This map was enacted and upheld by the Maryland Court of Appeals.
MD	Rep	Maryland Citizens Redistricting Commission	CRC Final Recommended Congressional Map	The governor traditionally proposes a map. The Republican Governor delegated this job to a task force called the Maryland Citizens Redistricting Commission.
NC	Dem	Senator Jay Chaudhuri	CST-6	Of the proposals submitted by Democrats, we chose the proposal that was later used in court cases challenging the Republican map.
NC	Rep	North Carolina General Assembly	Congressional Plan Senate Bill 740 / SL 2021-174	This proposal was enacted and later rejected by the state supreme court. However, the following year the state supreme court overturned their previous opinion.
NE	Dem	Democratic State Senator Justin Wayne	LB2	Although Nebraska has a nonpartisan legislature, this map was endorsed by the state Democratic party.
NE	Rep	Nebraska Legislature Redistricting Committee	LB1	This proposal was enacted and later challenged in state court.
NM	Dem	New Mexico Legislature	New Mexico Final Congressional Plan	This proposal was enacted and later challenged in state court.
NM	Rep	Citizen Redistricting Committee	Map A	This proposal was authored by the non-partisan redistricting committee but endorsed by the Republican party.
NV	Dem	Democrats in the Nevada Legislature	Nevada Final Congressional Plan	This proposal was enacted and later challenged in state court.
NV	Rep	Republicans in the Nevada Legislature	Congressional Minority Plan	This proposal was submitted by the Legislature Joint Minority Committee.
NY	Dem	Democrats in the New York Legislature	New York State Assembly Congressional Plan	The New York State Independent Redistricting Commission submitted two maps: one supported by Democrats and one supported by Republicans. The democratic legislature then proposed this map, which was enacted by the legislature and later rejected by the state supreme court.
NY	Rep	Republican-appointed members of the New York Independent Redistricting Commission	Plan B	This map was proposed by Republicans appointed to the New York State Independent Redistricting Commission.
PA	Dem	Senate Democratic Caucus	Senate Democratic Caucus	This plan was submitted by the Senate Democratic Caucus in the state supreme court case.
PA	Rep	General Assembly	Updated Preliminary Congressional Plan	This plan was enacted and was later vetoed by the democratic governor.

F.3. Geospatial analysis

We use geospatial analysis to merge precinct-level data in Python. The precinct boundaries vary slightly across the precinct-level data sources for election returns, party registration, and demographics. We aggregate all precinct-level data to the boundaries used by the 2020 Census. Where necessary, we use block-level population data to disaggregate precincts in the elections and voter registration data. Due to the discrepancies in precinct borders, we expect some measurement error in the final precinct-level dataset, especially when computing ratios like turnout rates in small precincts. We drop 4,633 precincts (8.3% of precincts) in which there is either a) a total population under 50, b) reportedly more votes than people, c) reportedly more registered voters than people, or d) reportedly more citizens of voting age than people. These tend to be small precincts, accounting for 2.3% of the total population in the sample.

Finally, we assign precincts to congressional districts for each map. We compute the percent of a precinct's land area that falls inside a district. A precinct is assigned to a unique district if at least 99.9% of its land area is inside one district. Otherwise, a precinct may be assigned to multiple districts (though we treat overlays that are smaller than 0.1% of the precinct's land area as error). Precincts are split infrequently: 1.4% of precincts are split in the Democrat proposals and 1.1% of precincts are split in the Republican proposals.

Where precincts span multiple districts, we disaggregate precinct characteristics using the percentage of precinct population in each district. To measure the population in a part of a precinct, we use block-level population data from the Census. In the rare case that a block is split across a precinct, we use the land area of the block to disaggregate the population of the block. We drop split precincts from the regression analysis in Section 6.2, since we use precincts as the unit of analysis and make comparisons across Democrat and Republican proposals. We do, however, include split precincts in the analysis of counterfactual maps in Section 6.3. There, we allow parts of split precincts to be swapped along the border when we construct counterfactual maps.

Our final precinct-level dataset includes 44,338 precincts across the 10 states (summary statistics are reported in Section F.5). For each precinct, we know the district it is assigned to under both proposals, allowing us to compute district-level characteristics. We also identify if a precinct is adjacent to a district border. We use this subset of border precincts to construct and evaluate counterfactual maps, as described in Section 6.3.

F.4. Measures of Turnout Rate

A key challenge when measuring turnout rate is that observed turnout is potentially endogenous to redistricting. In particular, turnout might become mobilized if a precinct is assigned to a competitive district and suppressed if it is assigned to a noncompetitive district.

We use three measures of turnout rates to address endogeneity concerns.

First, we use data from elections prior to redistricting to measure turnout. We compute the turnout rate of a precinct as the average of the total number of votes in the 2016 and 2020 presidential elections, divided by the population of the precinct in 2020.[32] In this way, election data do not come from contests directly affected by the 2020 redistricting proposals under consideration. However, the observed turnout might still be endogenous to redistricting proposals, if districts in the 2020 cycle are similar to districts in the 2010 cycle.

In a second approach described in the next section, we predict turnout using only demographic and socioeconomic characteristics, which are difficult for political parties to manipulate: population, race and ethnicity, citizenship, age, education, and income.

Just in case one worries that our prediction model is overly sophisticated compared to the information at hand for partisan gerrymanderers, in a third approach we use Citizen Voting Age Population (CVAP) as an exogenous proxy for the turnout rate (the coefficient of correlation between percent CVAP and turnout rate is 0.38).

F.4.1 Predicted Turnout

To predict turnout at the precinct level, we measure precinct-level demographic and socioeconomic factors using data from the American Community Survey and Census. The candidate variables and definitions are in Tables F.2 Population, race, and ethnicity variables are from the 2020 Census: P.L. 94-171 Redistricting Data Summary File (downloaded at the Census Voting Tabulation District level

[32] We take the average of the two presidential elections to limit noise. We use presidential elections instead of congressional elections or state-level elections because we do not need to impute values for uncontested elections and because within-state contest-effects, which could depend on redistricting, only indirectly affect presidential elections.

from NHGIS).[33] The Citizen Voting Age Population variables are from the Redistricting Data Hub.[34] The Redistricting Data Hub uses the 2019 CVAP special CVAP tabulation files from the American Community Survey (ACS) 5-Year Estimates (2016-2020). They disaggregate Census Block group-level estimates to the Census Block level. We then aggregate the block-level data to the precinct level, as in section F.1. Age, education, and income variables are from the 2020 American Community Survey: 5-Year Data (2016-2020). We download these data at the Census Tract level from NHGIS. To aggregate to the precinct-level, we first disaggregate to the block level (using percent of tract population in a block), then aggregate to the precinct level.

We compare the performance of three different models, namely Lasso, Random Forest, and Gradient Boosting. The dataset was divided into training and test data. The training data, comprising 75% of the precincts, was used to train each model, and the remaining 25% was used to evaluate accuracy.

The state was included as a categorical variable in all models. We additionally train a LASSO model separately for each state. We used the LassoCV function from scikit-learn, which automatically selects the alpha value that minimizes the Mean-squared error (MSE) by cross-validation. The maximum number of iterations was set to 100. For Gradient Boosting we use the HistGradientBoostingRegressor algorithm from scikit-learn. Table F.3 reports the MSE and the correlations between turnout and predicted turnout for the test data. We use predictions from the Gradient Boosting model in our main analysis since it has the lowest MSE. Figure F.1 plots the distribution of turnout and predicted turnout across all precincts. The most important predictors are reported in Table F.4, ranked by feature importance. Feature importance measures the decrease in model performance when the values of a variable are randomly reassigned. For each candidate variable, we randomly reassign values within the dataset, then train the model and measure MSE for the test data. We repeat this process 10 times for each variable. The feature importance score is the average difference in MSE caused by the permutation of the values of a given variable. Higher numbers indicate that the candidate variable has a larger effect on the model's performance.

[33] Steven Manson, Jonathan Schroeder, David Van Riper, Tracy Kugler, and Steven Ruggles. IPUMS National Historical Geographic Information System: Version 17.0 [dataset]. Minneapolis, MN: IPUMS. 2022. `http://doi.org/10.18128/D050.V17.0`.

[34] Available at: `https://redistrictingdatahub.org/`.

TABLE F.2. Candidate variables for turnout prediction models

Population, Race, and Ethnicity
- Total Population
- Voting Age Population (aged 18+)
- Total Hispanic or Latino
- Total White, non-Hispanic
- Total Black, non-Hispanic
- Total AI/AN, non-Hispanic
- Total Asian, non-Hispanic
- Total NHOPI, non-Hispanic
- Total Other Race, non-Hispanic
- Percent Hispanic or Latino
- Percent White, non-Hispanic
- Percent Black, non-Hispanic
- Percent AI/AN, non-Hispanic
- Percent Asian, non-Hispanic
- Percent NHOPI, non-Hispanic
- Percent Other Race, non-Hispanic

Citizen Voting Age Population (CVAP)
- CVAP, Total
- CVAP, American Indian or Alaska Native
- CVAP, Asian
- CVAP, Black or African American
- CVAP, Native Hawaiian and Other Pacific Islander
- CVAP, White
- CVAP, Hispanic or Latino

Age, Education, and Income
- Population aged 0 to 17
- Population aged 18 to 24
- Population aged 25 to 34
- Population aged 35 to 44
- Population aged 45 to 54
- Population aged 55 to 64
- Population aged 65 and up
- Percent aged 0 to 17
- Percent aged 18 to 24
- Percent aged 25 to 34
- Percent aged 35 to 44
- Percent aged 45 to 54
- Percent aged 55 to 64
- Percent aged 65 and up
- Population with no high school degree
- Population with high school degree
- Population with some college
- Population with Bachelor's degree
- Population with graduate degree
- Percent with no high school degree
- Percent with high school degree
- Percent with some college
- Percent with Bachelor's degree
- Percent with graduate degree
- Median household income
- Labor force population
- Percent of adults in labor force

Note: AI/AN is American Indian and Alaska Native, NHOPI is Native Hawaiian and other Pacific Islander.

TABLE F.3. TURNOUT PREDICTION: MODEL PERFORMANCE

Model	MSE	Correlation
Lasso, state-specific model	0.0100	0.6853
Lasso, pooled model	0.0126	0.5897
Random Forest	0.0065	0.8162
Gradient Boosting	0.0061	0.8270

TABLE F.4. Top predictors, ranked by Feature Importance

Variable	Feature Importance
Percent White, non-Hispanic	0.2446
State categorical variable	0.2213
Total population	0.0880
Percent with no high school degree	0.0490
Citizen Voting Age Population	0.0406
Percent below poverty rate	0.0406
Percent age 65 and up	0.0278
Percent with graduate degree	0.0237
Percent with high school degree	0.0222
Total Hispanic population	0.0217
Percent with Bachelor's degree	0.0184
Percent age 55 to 64	0.0173
Percent Black, non-Hispanic	0.0108
Percent Hispanic	0.0107
Percent of adults in labor force	0.0093

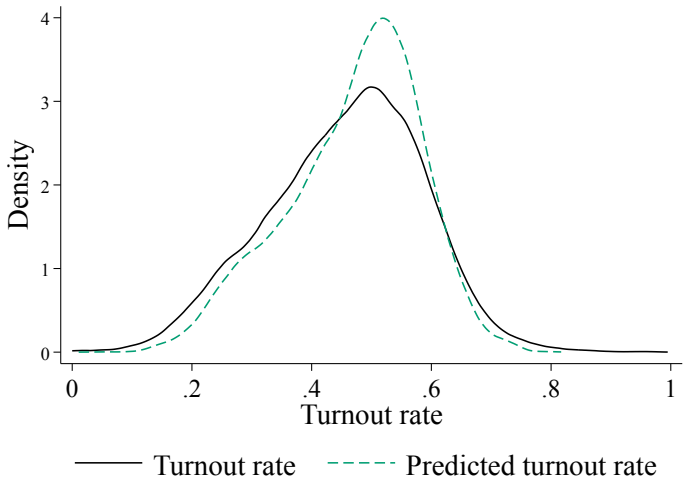


FIGURE F.1. DISTRIBUTION OF TURNOUT RATE AND PREDICTED TURNOUT RATE. The turnout rate is the total number of votes for both the Democrat and Republican parties divided by the total population. The black line is the average turnout rate in the 2016 and 2020 presidential elections. The green dashed line is the predicted turnout rate.

F.5. Summary Statistics

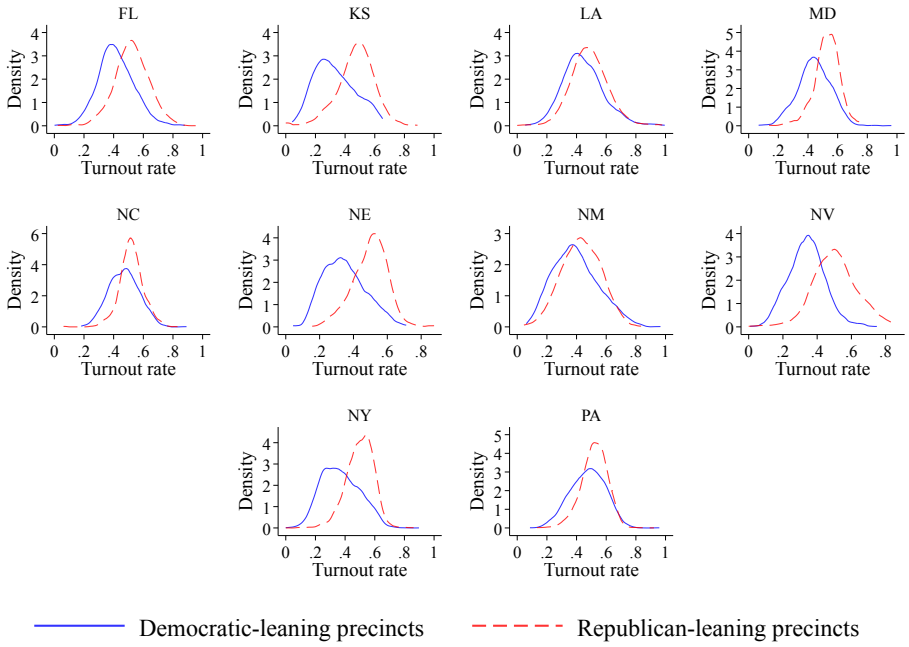
Table F.5 reports summary statistics for all precincts. Figure F.2 shows the distribution of turnout rates, by partisan lean, for each state.

TABLE F.5. Summary statistics for precincts (N=43,390)

	Mean	SD	p1	p25	p50	p75	p99
Population, Eligibility, and Turnout							
Total population	1,919	1,731	89	890	1,429	2,340	8,435
Percent voting age population	.79	.06	.64	.76	.79	.82	.96
Percent citizen voting age population	.71	.14	.35	.63	.72	.81	.98
Percent registered voters	.64	.14	.23	.56	.65	.73	.92
Turnout rate	.46	.13	.15	.37	.47	.55	.74
Partisan Affiliation							
Percent registered Democrat	.57	.24	.11	.38	.55	.78	.98
Percent registered Republican	.43	.24	.023	.22	.45	.62	.89

Note: Mean, Standard Deviation (SD) and percentiles (p) for precincts in the following 10 states: FL, KS, LA, MD, NC, NE, NM, NV, NY, PA.

FIGURE F.2. Distribution of turnout rate across precincts, by partisan lean



Note: This figure shows the distribution of turnout rate across precincts of a given partisan lean for each state. A precinct leans Democratic if it has an above median share of registered Democrats, otherwise it leans Republican.

G. Additional Empirical Evidence

G.1. Gerrymanderers Take Turnout into Account

Figure G.1 shows the number of expected Democrat seats under each proposal. In 8 of the 10 states, Democrats expect to win strictly more seats under the Democrat proposal than the Republican proposal. In total, if all Democrat proposals were implemented there would be 14 more seats for Democrats than if all Republican proposals were implemented, a 12% swing in seats. Notably, if we were to use registered Democrats and Republicans to evaluate these maps, ignoring turnout data, then there would only be 7 seats of difference between Democrat and Republican proposals, a 6% swing in seats.

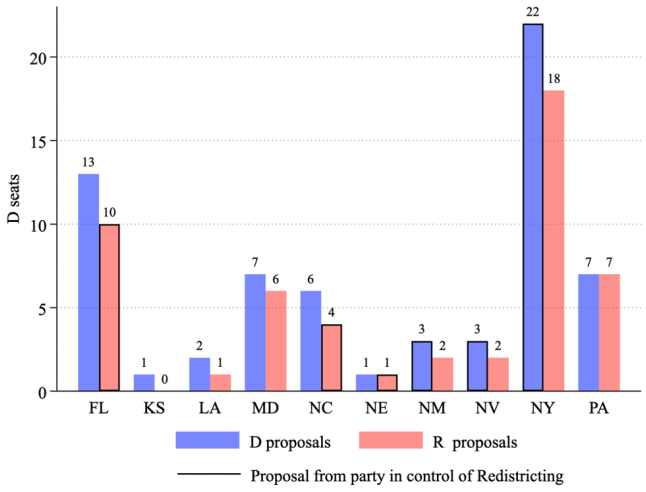


FIGURE G.1. NUMBER OF EXPECTED DEMOCRAT SEATS UNDER BOTH PARTIES' PROPOSALS FOR U.S. CONGRESSIONAL DISTRICTS. This figure shows the expected number of seats won by Democrats under the Democrat (D) and Republican (R) proposals for U.S. Congressional districts. A seat is expected to be won by Democrats if the Democrat vote share is larger than the Republican vote share. Vote shares are measured using the average of the two presidential elections prior to the redistricting cycle (2016 and 2020). The party in control of redistricting is outlined in black. A party is in control of redistricting if both the majority party in the legislature and the party of the governor are aligned. Three states have split control: LA, MD, and PA.

G.2. Alternative definitions of partisan lean

This section evaluates the robustness of results in Section 6.2 to alternative measures of partisan-lean. Recall that we find a negative correlation between a precinct turnout rate and the change in Democratic vote share of a district under the Democrat vs. Republican proposal, within precincts of a given partisan lean. In Table 1, we exclude precincts for which there is no obvious partisan lean. These so-called weakly partisan precincts have a two-party share of registered voters between the 25th and 75th percentile of percent Democrats (between 38% and 78% Democrat). Figure G.2 shows the distribution of percent Democrat across all precincts. The distribution is bi-modal and the median value is 0.55.

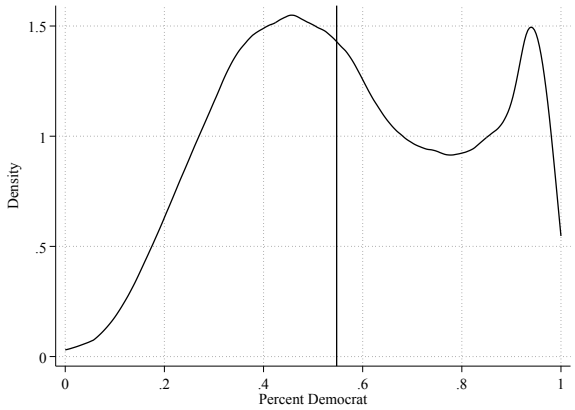


FIGURE G.2. DISTRIBUTION OF PERCENT REGISTERED DEMOCRATS WITHIN PRECINCTS. Percent Democrat is the two-party share of registered voters that are Democrats. The vertical black line indicates the median value of 0.55.

In Figure G.4, we show OLS regression coefficients for estimating Equation 5 using alternative percentiles to define weak precincts. As the percentile threshold increases, only more strongly Democrat- and Republican-leaning precincts are included in the sub-sample. The negative correlations in Table 1 are robust to these alternative thresholds. While the coefficients are small in magnitude and statistically insignificant for the full sample, the magnitudes of the coefficients increase as the threshold for partisan lean increases and sub-samples include more strongly partisan precincts.

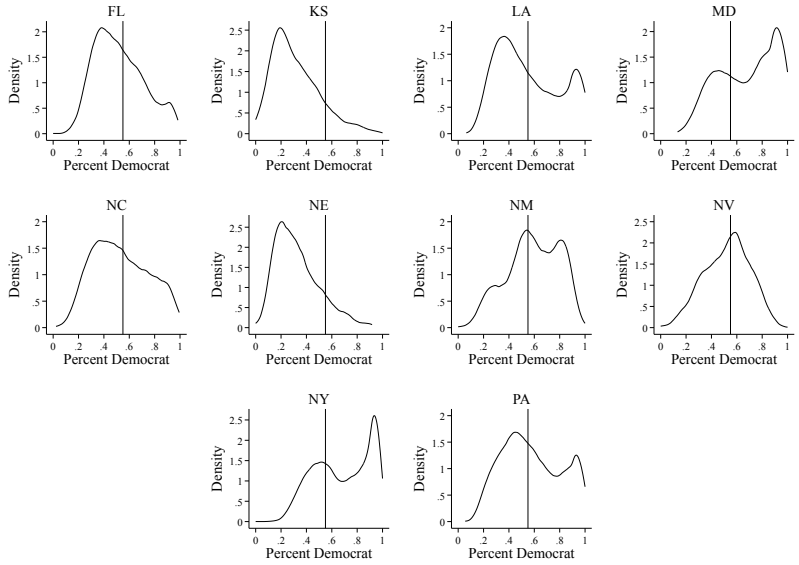


FIGURE G.3. DISTRIBUTION OF PERCENT REGISTERED DEMOCRATS ACROSS PRECINCTS, BY STATE. It shows the two-party share of registered voters that are Democrats. A vertical black line indicates the median for the full sample: 0.55.

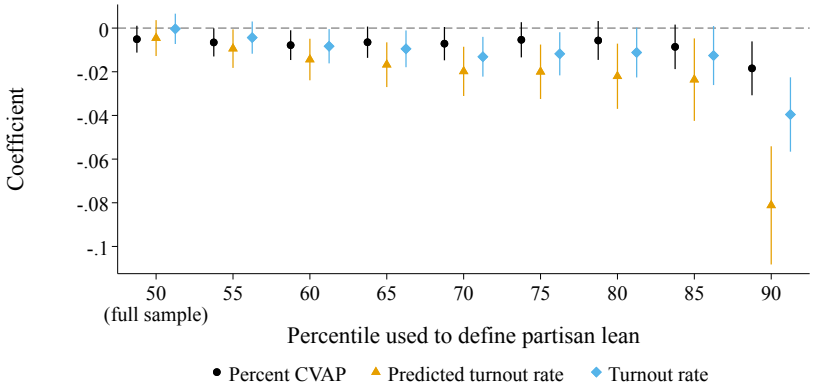


FIGURE G.4. CORRELATION BETWEEN TURNOUT AND CHANGE IN DEMOCRATIC VOTE SHARE BY STRENGTH OF PARTISAN LEAN. This figure plots the regression coefficients from estimating Equation (5) for sub-samples of precincts based on the definition of partisan-lean. Each symbol is a point estimate from a separate regression in a sub-sample of precincts in which the two-party share of Democrats or Republicans exceeds the percentile indicated on the x-axis. Vertical lines indicate 95% confidence intervals. The independent variable is one of three measures of turnout: percent Citizen Voting Age Population (CVAP), predicted turnout rate, or turnout rate.

G.3. Calibrating standard deviation of the aggregate shock

We use the standard deviation of two-party Democratic vote share from recent presidential elections (2008-2020) to calibrate σ_s, the standard deviation of the aggregate shock in a state.

TABLE G.1. STANDARD DEVIATION OF STATEWIDE DEMOCRAT VOTE SHARES, 2008-2020

State	σ_s
Florida	0.013
Kansas	0.021
Louisiana	0.006
Maryland	0.019
Nebraska	0.025
Nevada	0.024
New Mexico	0.013
New York	0.021
North Carolina	0.009
Pennsylvania	0.025
Average	0.017

G.4. Swapping Precincts at the Border: Robustness Checks

Most swaps in Figure 2 of Section 6.3 have a negligible effect on the gerrymanderer's payoff function. In Figure G.5, we plot the percent of swaps that are profitable for Democrats and Republicans, varying the definition used to define a profitable swap. That is, we say a swap is profitable only if the effect on the payoff is sufficiently high. The pattern in Figure G.1 holds if we ignore swaps with a negligible effect on the payoff: each party has fewer unexploited profitable swaps under their own proposal than under their opponents' proposal (Figure G.5, Panel A). For example, 5% of swaps are profitable for Democrats by 1 p.p. or more under their own proposal, versus 8% under the Republican proposal. Only 4% of swaps are profitable for Republicans by 1 p.p. or more under their own proposal, versus 6% under the Democrat proposal.

The pattern eventually reverses, however, if we focus only on the most extremely profitable swaps. Democrats have more profitable swaps with an effect of 2.5 percentage points or more under their own proposal than under the Republican proposal. Note that the swaps with the largest effects might be

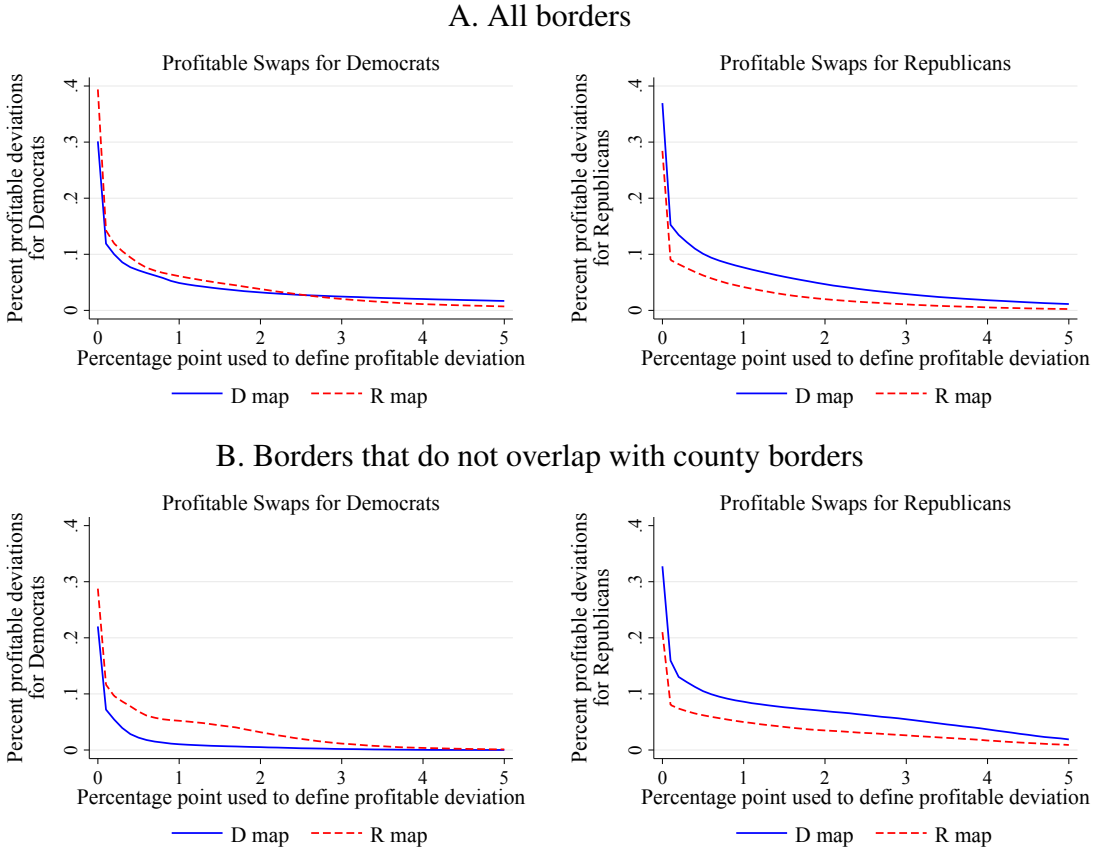


FIGURE G.5. SHARE OF SIGNIFICANTLY PROFITABLE SWAPS. The y-axis is the percent of feasible counterfactual maps ('swaps') that are profitable for Democrats (left) and Republicans (right). A swap is profitable if the change in payoff exceeds the threshold on the x-axis. The payoff is the expected number of seats. To compute the expected number of seats we assume that the aggregate shock is normally distributed with mean zero and standard deviation calibrated using statewide presidential election returns from 2008-2020 (see Appendix G.3, values range from 0.01 to 0.02).

especially difficult to implement due to legal and geographic constraints. For example, counties should not be divided into multiple districts, where possible. At the same time, there is more variation in vote shares across counties than within. Thus, where a district border coincides with a county border, swaps are more likely to have a significant effect and less likely to be feasible. If we exclude borders that coincide with counties from our analysis, we see that parties have fewer profitable deviations overall and that Empirical Prediction 2 holds, even for swaps with large effects on payoffs (Figure G.5).

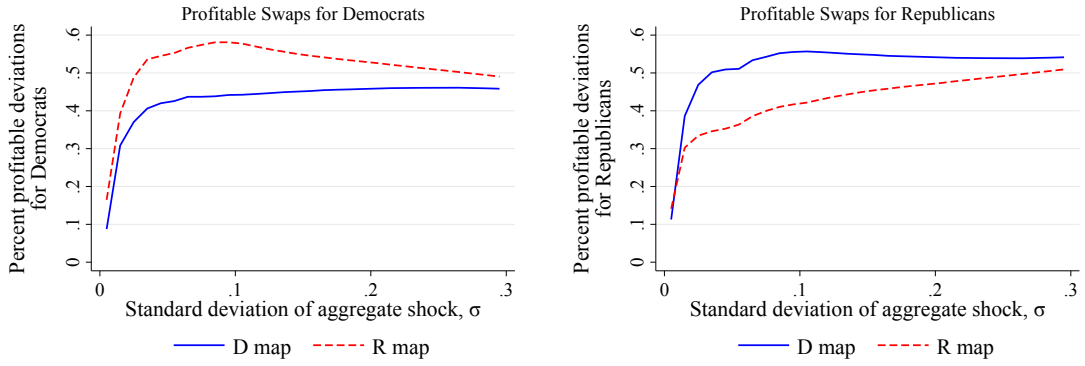


FIGURE G.6. SENSITIVITY ANALYSIS: profitable swaps by value of standard deviation of the aggregate shock (σ)

Next, we show that the pattern in Figure G.1 is robust to alternative values of σ, the standard deviation of the zero-mean normal distribution for the aggregate shock. We plot the share of profitable deviations by σ in Figure G.6. While the data-implied values of σ lie between 0.01 and 0.02, depending on the state, the difference in the share of profitable deviations across proposals at first increases for larger values of σ, peaking at 0.09. For large values of σ, the proposals tend to look more similar in terms of share of profitable deviations, approximating the case of the uniform aggregate shock.

Finally, we repeat these sensitivity analyses for the alternative methods of measuring payoffs using scaled votes and scaled registered voters in Figures G.7 and G.8.

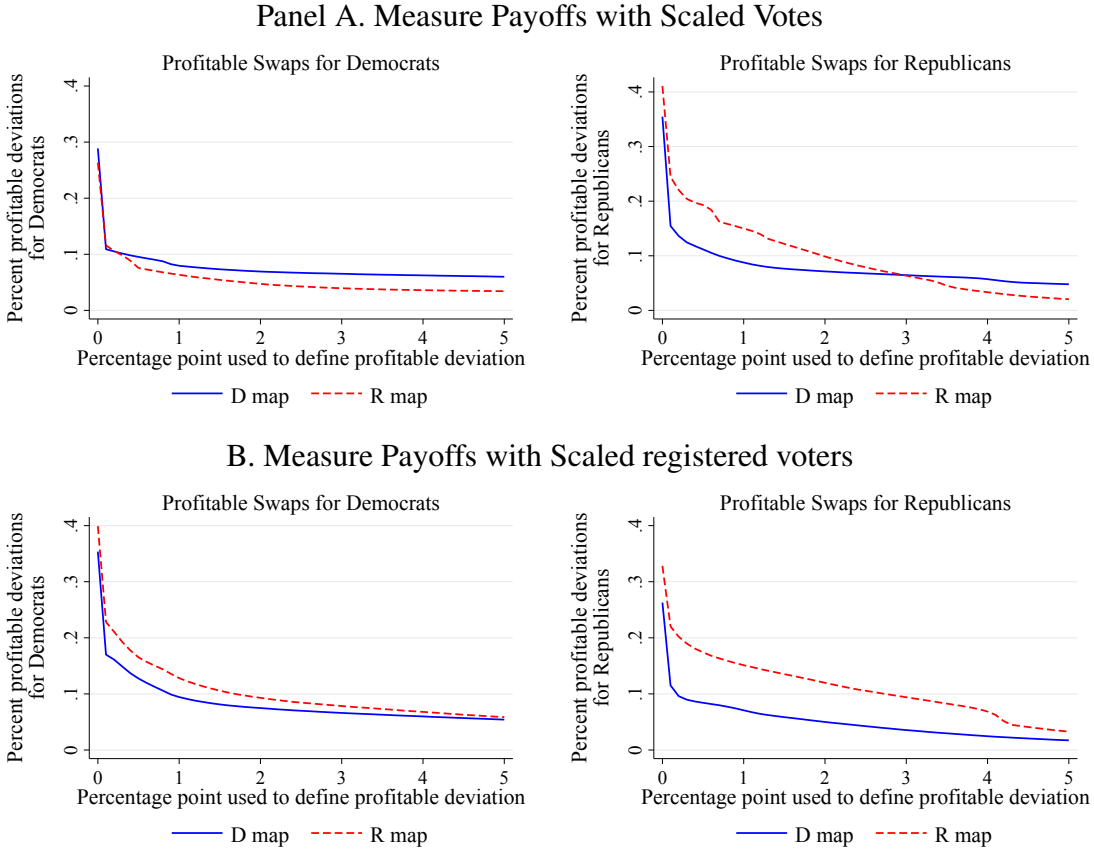


FIGURE G.7. SHARE OF SIGNIFICANTLY PROFITABLE SWAPS. The y-axis is the percent of feasible counterfactual maps ('swaps') that are profitable for Democrats (left) and Republicans (right). A swap is profitable if the change in payoff exceeds the threshold on the x-axis. The payoff is the expected number of seats. To compute the expected number of seats we assume that the aggregate shock is normally distributed with mean zero and standard deviation calibrated using statewide presidential election returns from 2008-2020 (see Appendix G.3, values range from 0.01 to 0.02).

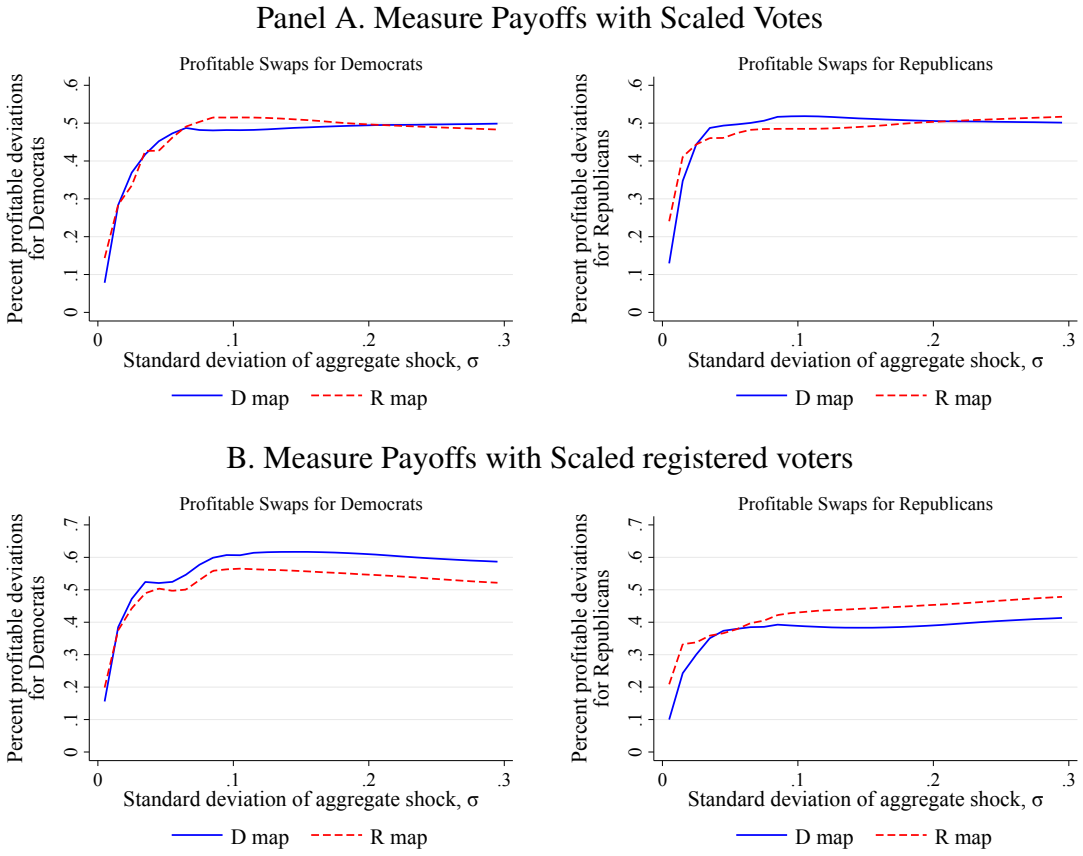


FIGURE G.8. SENSITIVITY ANALYSIS: profitable swaps by value of standard deviation of the aggregate shock (σ)