American Economic Journal: Economic Policy 2022, 14(4): 1–25
https://doi.org/10.1257/pol.20190213

AQ1

Gerrymandering in State Legislatures: Frictions from Axiomatic Bargaining†

By Hisam Sabouni and Cameron A. Shelton*

Theories of partisan redistricting postulate unitary actors maximizing their party's expected seat share. Yet, the partition of a fixed supply of friendly voters necessarily implies a tragedy of the commons. We recast partisan redistricting as a bargaining game among the sitting representatives of the party controlling the map. The status quo is the threat point, explaining why changes are frequently minor. This bargaining framework implies that highly competitive districts will receive more help from redistricting if they are already represented by the party in charge. Employing a regression discontinuity design with precinct-level data, we find support for this prediction. (*JEL* C78, D72)

At least since the work of Douglas North, economists have noted that political institutions play a critical role in economic development. Among others, Acemoglu, Johnson, and Robinson (2005) have argued persuasively that a society may fail to undertake productive investments and reforms if the distribution of economic returns is not congruent with the distribution of political power. Much of the related literature focuses on explaining the divergence between countries that have successfully industrialized and those that remain near subsistence income. However, a similar logic—whereby the concentration of political power can block efficient reforms—applies to the institutions of developed countries. Aghion, Alesina, and Trebbi (2004) note the importance of electoral institutions in determining the minimum size of the minority required to pass or block reforms. For instance, according to 2019 Census Bureau estimates, the 21 least populous states, controlling sufficient seats in the US Senate to block legislation via filibuster, contain just 11 percent of the US population. To the extent that these states and senators identify common interests—and the dominant policy questions are, at present, increasingly associated with the relevant urban–rural cleavage (Rodden 2019)—this constitutes a minority capable of blocking reform as a result of US electoral law. While, in the US Senate, the advantage of certain regions is permanently ensconced in the constitution, most

* Sabouni: Department of Economic Sciences, Claremont Graduate University (email: hsabouni@cgu.edu); Shelton: Robert Day School of Economics and Finance, Claremont McKenna College (email: cshelton@cmc.edu).
Dan Silverman was coeditor for this article. Special thanks to three anonymous reviewers. All errors are our own.

† Go to https://doi.org/10.1257/pol.20190213 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

other electoral districts in the United States are subject to decadal redistricting, frequently under the control of elected officials.

Partisan gerrymandering[1] is an attempt by a minority to perpetuate a temporary electoral success. Prior literature has noted that it leads to a bias of the seats–votes curve[2] in favor of the party currently in power (Coate and Knight 2007), leading to a legislature overrepresenting one set of interests. In addition to the direct effect on policy through party seat shares, Besley and Preston (2007) note that this also leads parties to shift their platforms in the direction of the bias. By solidifying control by whichever party happens to be in power at the time of redistricting, gerrymandering selects which minority rules—and, thus, which reforms—are passed as well as the distribution of net benefits of those reforms. Moreover, as we will show, gerrymandering ensconces not just which party rules, but which geographic districts are represented by that party. This is relevant given the myriad examples of district-specific benefits accruing to representatives in the majority (e.g., Barry, Burden, and Howell 2010), the effects of which can endure for decades (Ejdemyr, Nall, and O'Keefe 2015).

While attention paid to state redistricting has waxed in recent years, it is not clear whether we have the proper positive model to explain the districts that emerge under the present system. Theories of partisan gerrymandering based on explicit optimization typically assume a unitary actor free from the constraints of history. Some individual is presumed to draw the lines subject only to the constraint of equal population and to a degree of uncertainty over the future partisan preferences of voters.[3]

Yet, parties are not unitary. Members face a complex combination of individual and collective goals. The partitioning of a fixed electorate with limited copartisans implies a dilemma for any individual member. Maximizing expected party seat share entails transferring friendly voters from safe districts, where their marginal effect on probability of victory is low, to competitive districts where their marginal effect is high. Members in relatively safe districts are thus cross pressured. Improving the party's expected seat share may require that they reduce their own probability of reelection. Even members who would receive an influx of friendly voters under the optimal-seat-share scheme would likely prefer more help than the optimal scheme would allocate them. It is a common pool problem from which the map that maximizes expected seat share is unlikely to emerge.

We introduce a new model of gerrymandering that reflects this common pool problem. We take as our starting point the model of Gul and Pesendorfer (2010), which postulates that mappers observe vote choice from which they must infer voters' true preferences. Mappers are then constrained to rearrange blocks of

[1] Partisan gerrymandering is the drawing of districts in such a way as to maximize the likelihood of or size of a majority for the party controlling the process. Bipartisan gerrymandering—when neither party has a monopoly on control—purportedly results in the drawing of districts in such a way as to protect incumbents of both parties.

[2] The seats–votes curve is the function translating a party's vote share into its seat share. It is biased toward a party when it passes through (50, 50+x) where a party that receives 50 percent of the vote receives 50+x% of the seats for $x > 0$. The steepness of the curve is called the responsiveness, tracking how swiftly seat share rises in response to gains in vote share, and is typically in excess of 1.

[3] In an exception that does investigate an additional constraint, Fryer and Holden (2011) show that one particularly common legal constraint—the principle that mappers should attempt to keep districts compact—has an effect on the seats–votes curve. They find that compactness raises responsiveness but has no effect on bias.

voters whose average partisan composition is inferred from prior voting data. This approach has the advantage of being both microfounded and easily brought to available precinct-level data. But instead of presuming that the map is designed by a unitary actor to maximize expected seat share, we postulate that the map is the product of a bargaining process among the sitting legislators.

While most states pass the new map through standard legislative procedure, the majority party typically introduces a map that has already been agreed upon amongst its members. Written accounts of the intraparty bargaining process are unavailable, making it challenging to write a structural model of the relevant bargaining game. Thus, we choose an axiomatic approach to bargaining, the Nash bargaining solution, in which all seated legislators of the party in power are represented in the bargain. Legislators balance two considerations: they wish to maximize their own chance of reelection, and they wish to maximize the seat share of their party. From a holistic perspective, districts that are already held are represented in the bargain through both terms; they count toward party seat share, and one particular actor in the bargain—the current holder of the seat—places extra weight on the party winning that particular seat.[4] By contrast, districts held by the other party are part of the expected seat-share calculation but do not have a specific advocate at the bargaining table. As a result, friendly voters are, compared to the seat-share-optimizing level, overly allocated to seats already held, resulting in a departure from the optimal map.

Bargaining outcomes famously depend on the allocation that would prevail in the event that no agreement is struck. Cox and Katz (2002) convincingly argue the importance of the reversionary outcome in the context of redistricting. Most state constitutions set a deadline for drawing the lines. If the legislature cannot agree on a map before the deadline, the state might either use a backup commission or remand the issue to the state court. In the 2000 redistricting cycle, courts drew the lines in 11 of the 50 states (Levitt 2010, 28). In some cases, these backstops will share the partisan orientation of the legislature itself, but in most cases, the alternative is less partisan. As Levitt puts it, "Judges have little direct stake in the contours of particular legislative district lines, and may appoint individuals who similarly have little direct stake in the outcome of the redistricting process" (Levitt 2010, 28). To capture this lesser partisanship and the sense that caretaker commissions with little time and expertise are both less likely to radically redraw the map and likely to revise along lines orthogonal to partisan considerations, we presume that the disagreement point is the status quo.

The Nash bargaining solution maximizes the product of each player's surplus beyond the disagreement point. This results in relatively few changes being made from the previous map, as any change must deliver a Pareto improvement among the majority-party representatives. Thus, while existing unitary actor theories are entirely without inertia, results from our bargaining procedure are history dependent, with the previous map acting as an important reference point. We believe that this is important to explain what seems to be a high degree of inertia in the lines.

[4] We have clearly abstracted from considerations of primary challenges.

One of the axioms of the Nash solution is symmetry, which could be questioned on the grounds that more senior members of the chamber might have better connections and more accumulated favors to call in and, thus, achieve a more favorable position in the bargain. In response, we note that more senior members are also more likely to be in relatively safe seats, from which losing friendly voters is less costly. As a result, if one adds seniority to the model, the predicted relationship between seniority and voter allocation is of indeterminate sign, and the two effects cannot be distinguished empirically. Thus, we keep to the simple version for clarity.

In actuality, the perfect gerrymander is probably not possible, because the geographic distribution of voters, combined with traditional districting principles (TDPs) such as prioritizing compactness, prevents certain combinations of voters (Sabouni and Shelton 2020). It has been noted that Democrats naturally pack themselves into cities (Chen and Rodden 2013), making it easier for Republican gerrymanderers to waste Democratic votes without drawing salamander-esque shapes that could run afoul of compactness criteria. It has also been noted that the Voting Rights Act of 1965 provides natural Republican gerrymandering by requiring concentrations of minority voters who tend to vote Democratic (Shotts 2001). Both cases suggest a structural bias toward one party. The shape of a state and the geographic distribution of partisan affiliation may provide subtler bias in either direction. Indeed, recent lawsuits by reformers rest on Markov chain Monte Carlo explorations of the set of valid maps so as to enable comparisons of the chosen map with the distribution of potential maps and thereby investigate the claim that "geography made me do it." (Cho and Liu 2016).

The presence of geographic constraints would not interrupt our results. Theoretically, once one has identified which districts are to be packed, they are pushed to the feasible limit, yielding a maximal number of net friendly voters over which the remaining districts compete. The resulting set of feasible bargains remains convex in utility space, thus ensuring that the Nash bargaining solution remains valid. Empirically, existing districts are likely heterogeneous in their geographically determined potential for adding or donating friendly voters. This unobserved constraint could conceivably be correlated with the existing partisan balance of a district. If copartisans tend to cluster, then districts that straddle the boundary between opposing clusters are likely to be both closely contested and more easily adjusted. However, our regression discontinuity design is sufficient to address precisely this situation.

The Nash bargaining solution delivers results that are intuitive yet offer new insight. We reaffirm that the party in charge will collectively reassign friendly voters from safe districts to competitive districts. This acts to rebalance the electorate toward the optimal crack-and-pack solution. However, there are two important differences. First, this process is limited by the necessity of delivering any particular representative a positive net utility. Thus, even the safest district currently held by the party will become only modestly more competitive before the holder of that district finds that donating additional voters outweighs the internalized benefits of a larger majority. Second, the process privileges districts that are already held by the majority party and, thus, represented in the bargaining process. Suppose there exist in the chamber two seats whose electorates each appear to be a knife-edge

50-50 split. In the most recent election prior to redrawing the maps, one of them just tipped for the opposition, while the other just tipped for the party in power. From the perspective of maximizing seat share, they should both receive equivalent assistance.[5] But because the latter is represented in the bargain, that district's changes will be more favorable. Our model thus delivers an important testable implication that enables us to evaluate the suitability of the unitary actor assumption.

Using precinct-level data, we calculate the change in normal Democratic vote that results for each district from the 2010 redistricting wave. This requires geographic information system (GIS) work with precinct boundaries and district boundaries to figure out which precincts are reallocated. The change in a district's normal Democratic vote represents the degree to which the majority-party mapmakers are either improving or reducing the chances of their candidate's victory. Restricting to those chambers with clear partisan control of the redistricting process and using previous vote margin as the running variable, we conduct a regression discontinuity analysis to measure the bias toward districts currently held by the party in power. The results support our theory in several ways. First, there is a statistically significant bias toward currently held districts of approximately 2 percentage points. Second, this bias is larger in chambers that are uncompetitive, where we would expect individual reelection considerations to trump the diminished value of expanding an already large seat-share majority. Third, friendly voters are transferred on net from safe seats to competitive seats.

In short, there is renewed interest in redistricting including the degree to which partisan gerrymandering inhibits fair representation and the conditions under which such behavior is constrained. To answer these questions requires an accurate model of the process. We believe that the unitary actor assumption upon which existing theories are based exaggerates the ability of the majority party to coordinate to solve the common pool problem. We build a model based on bargaining among existing members of the majority-party caucus, which delivers a testable hypothesis for which we find support in data from state legislatures.

The remainder of our paper consists of six sections. First, we present an example to give a clearer picture of the seemingly missed opportunities that we see as we scrutinize district maps and that motivated this theory. Second, we briefly review the two most prominent prior theories of gerrymandering. Third, we develop the model of mapmaking as a bargain. Fourth, we present our data. Fifth, we conduct our empirical work. Finally, we conclude with discussion of the implications and ideas for further work.

I. Example: Michigan State Senate Districts 9 and 10

We present an example chosen from a state with unified control of the process by a single party and without preclearance requirements.

[5] Incumbent advantages would break the symmetry. The majority party would enjoy a higher probability of victory in the seat currently held and a lower probability of victory in the seat held by the opposition. So long as both seats remain within the set to be helped by redistricting (and given that they are the 50-50 seats, they will so remain), this means that the seat held by the opposition would, in a unitary model, receive more assistance rather than less. Thus, the results of our regression discontinuity cannot be explained by incumbent advantage.

In the redistricting wave following the 2010 census, Republicans enjoyed unified control over the redistricting process in Michigan. Each state senate district comprised roughly 260,000 people, and the Republican Party controlled 26 of the 38 seats following the 2010 election. District 9, centered on the northern Detroit suburbs of Warren and St. Claire Shores, had been lost to the Democrats by a mere 9,000 votes in a relatively close election. Its neighbor to the north, District 10, had been won by a similarly close margin of 7,000 votes. District 9's neighbor to the north, District 11, had been won by a more comfortable margin of 31,000 votes (see Figure 1).[6]

A feasible local goal for a unitary actor would be to flip District 9 and secure District 10 via exchanges with Districts 2, 11, 12, and 13. The strategy to a mapper with our data on normal Democratic vote share by precinct is pretty clear. There is no way to flip District 2, so one ought to further pack it via exchange with District 9. Specifically, take the more conservative lakeside Grosse Point neighborhood from District 2 in exchange for some of the bluer inland bits of District 9. Secondly, one ought to secure District 10 by exchanging voters with 11 (and possibly 12 and 13). For instance, widen that liberal foot of District 11 to include more of District 10's local blue enclave in exchange for inland voters from District 11. Neither geography nor the local availability of copartisans should inhibit this seat gain.

However, in the event, the mappers did not pursue this course (see Figure 2). District 11 did send a large square of conservative precincts to District 10 but did not take any liberal voters in return. Instead, District 11 received the redder lakeside neighborhoods of District 9 by way of compensation, thus not materially altering the partisan makeup of the district. Meanwhile, District 10 offloaded many of its bluest precincts to District 9. In short, District 9 was written off and deliberately packed to shore up District 10, and Districts 11, 12, and 13 made no sacrifices despite having won by comfortable margins. In the first election under the new map, District 10 was retained by 21,000 votes, while District 9 was lost by 25,000 votes. Districts 11, 12, and 13 were all comfortably retained. It would seem that Districts 11, 12, and 13 were unwilling to shade their margins of victory in an attempt to flip District 9.

We are not claiming to have proven this to be an error—it is possible that this was the proper choice to maximize expected seat share given expected vote share volatility over the next decade. Our argument rests on the statistical case to be made later. We simply offer this as an example of what such missed opportunities look like. Our fundamental point is that the assumption of maximization of expected seat share that underlies the current models of gerrymandering does not seem to do well in describing the maps produced.

II. Prior Models

There are two prominent prior theories of partisan gerrymandering. The first is the well-known crack-and-pack theory (Owen and Groffman 1988; Gul and Pesendorfer 2010). This theory predicts that the party drawing the lines will pack as many

[6] Michigan State Senate election numbers are drawn from Klarner (2018). Boundary files come from the US census.

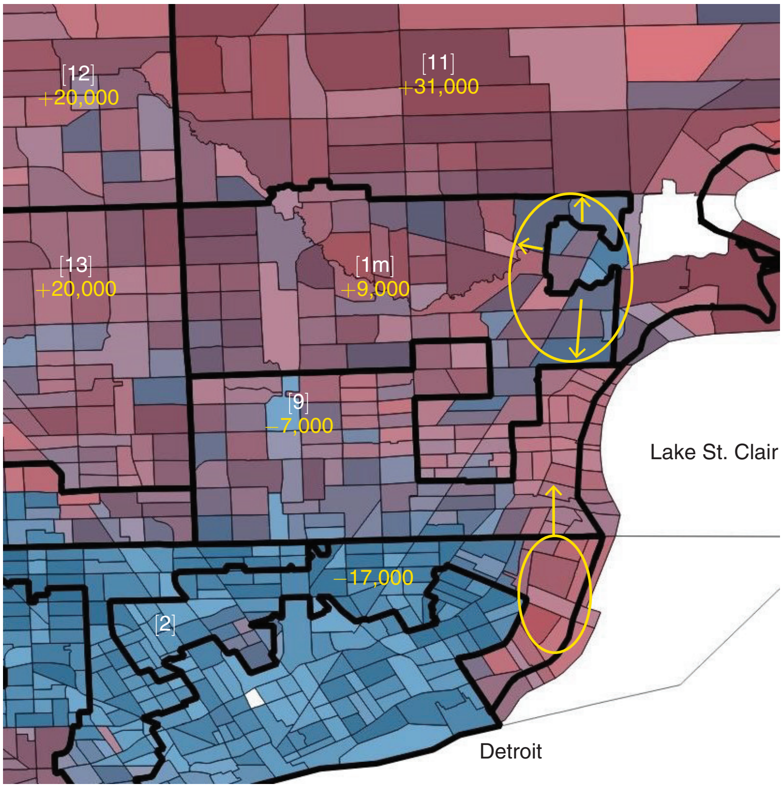

Figure 1. Michigan State Senate District Map for the 2000s

Notes: The smallest geographic units are voting precincts. The darker outlines are the state senate districts, indicated with numeric labels in brackets. The change in hue from red to blue indicates variation in the normal democratic vote from Republican to Democrat. The shading from light to dark indicates population. The yellow numerals indicate the electoral margins of the mapping party in the most recent election. The yellow arrows are suggestions developed in the text.

opposition voters as possible into a few districts, enabling the mapmaker to create a large number of districts with modest yet secure majorities. From a wasted-votes perspective (McGhee 2014), the excessive majorities in opposition districts waste far more votes than the modest majorities in supportive districts, allowing the mapmaker's party to achieve seat share in excess of vote share. Crack-and-pack predicts large majorities in opposition districts and smaller majorities in government districts. It also predicts that voters in opposition districts will be relatively similar, while voters in government districts will be more widely dispersed along the political spectrum.

By contrast, Friedman and Holden (2008) suggest that optimizing mapmakers will slice the partisan distribution of voters and match from the outside in. That is, a slice of the most implacable foes will be matched with a modestly larger slice of the most stalwart supporters in a single district. The mapmaker will then slice the remaining voters in the same way, covering the most ardent foes with a slightly larger set of the most ardent supporters. And so on, until the final district is a single

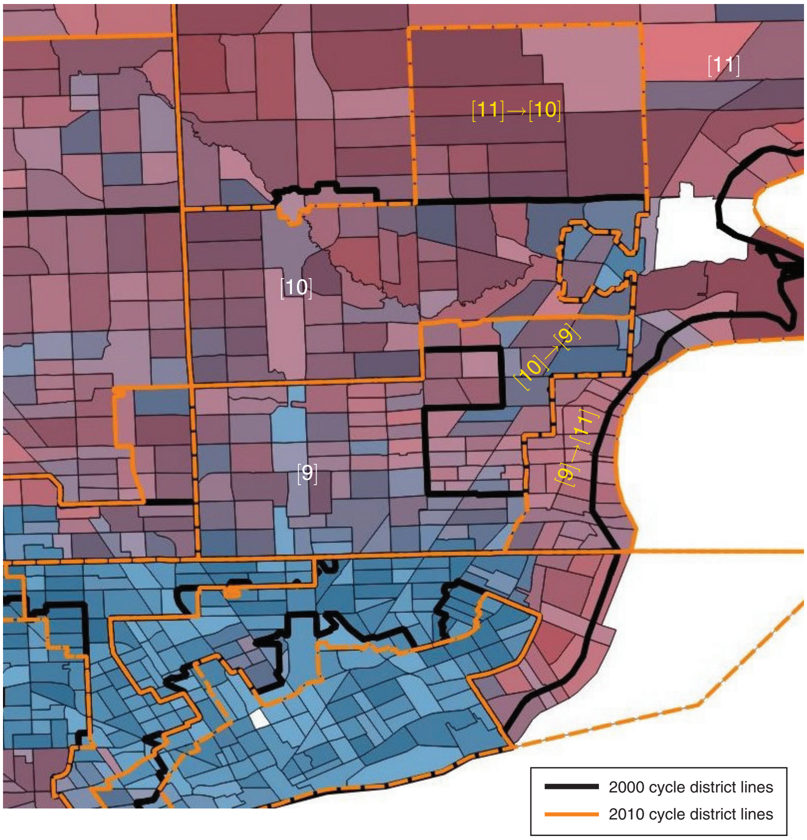

Figure 2. Michigan State Senate Redistricting for the 2010s

Notes: The smallest geographic units are voting precincts. The darker outlines are the state senate districts, indicated with numeric labels in brackets. The change in hue from red to blue indicates variation in the normal democratic vote from Republican to Democrat. The shading from light to dark indicates population. The yellow text labels "X→Y" indicate areas that are reallocated from district X to district Y.

lump of voters from near the center of the distribution. In each case, the slice of supporters is somewhat larger than the slice of foes so as to ensure that the median is a friendly voter. The size of this "overbite" depends on the degree of volatility in voter preferences. Friedman and Holden (2008) show that, in the face of uncertainty over voters' true preferences, this "slice-and-cover" improves the likelihood that districts set up to favor the mapmaker's party will actually be won by the mapmaker's party after preference-uncertainty is resolved.

Both Gul and Pesendorfer (2010) and Friedman and Holden (2008) offer compelling microfoundations in the form of voters with latent preferences about which parties receive signals. We have chosen to base our theory on the former partly because our analysis of maps and reading of commentary suggests that crack-and-pack behavior is more widespread than slice-and-cover, but also because the structure is more easily matched to existing data. To this, we have added the bargaining element that delivers the common pool problem and the testable hypotheses.

III. Theory

A. *The Model*

In this section, we first adopt and briefly describe the stochastic median voter model of Gul and Pesendorfer (2010). Parties are presumed to have fixed positions. Voters have symmetric, single-peaked preferences over a unidimensional policy space admitting an ideal point, x. While the ideal point is unknown to the mapper, the mapper receives signals on the partisan affiliation of a voter. Republicans have ideal points drawn from the cumulative distribution I_R, while Democrats have ideal points drawn from I_D. Republican policy is fixed at $+1$, while Democratic policy is fixed at -1. Thus, a Republican voter is one whose ideal point is greater than zero, while a Democratic voter has an ideal point less than zero: $I_R(0) = 0; I_D(0) = 1$. It is assumed that I_R is strictly increasing and convex on $[0,1]$, has a median in $[0,1]$, and is continuous. Symmetric conditions hold for I_D.

We draw attention to three substantive assumptions inherent to this setup. First, the median Republican voter is more moderate than the Republican platform (and likewise for Democrats). Second, the convexity assumption—which is crucial for a later result—implies that the density of Republican voters increases along the interval $[0,1]$, which implies the distribution of voters within a party is unimodal but the distribution of all voters across both parties is bimodal. Third, voters are classified according to whichever platform is closest to their ideal point. As this is a stochastic voting model, that correctly tracks the party the voter will more commonly vote for. This accords with our focus—which Gul and Pesendorfer (2010) seem to share—on the information derived from precinct-level voting records rather than measures of actual party membership.

Voters' utility depends on the distance between their ideal point and the platform of the party whose candidate is elected plus a valance term, v, drawn from cumulative distribution $L(\cdot)$. Positive values of v are presumed to favor the Democratic candidate. Thus, a voter with ideal point x receives utility $u_R(x,v) = -|1-x| - v$ if the Republican candidate is elected and $u_D(x,v) = -|-1-x| + v$ if the Democratic candidate is elected. The voter will prefer the Republican candidate if and only if $v < x$ for $x \in [-1,1]$. We assume L is strictly concave on $\mathbb{R}_+$, continuous, and symmetric around 0.

The chamber consists of N districts, currently split between the two parties, R and D, controlling N_R and N_D districts respectively. Without loss of generality, assume $N_R > N_D$ such that Republicans are the majority party in control of the redistricting process. By observing the precinct-level vote shares of each party, the mapper can construct the fraction of voters in each district that are affiliated with each party.[7] Prior to the current round of redistricting, each district i is endowed with a set of voters characterized by a fraction p_i affiliated with party R (thus $1 - p_i$ affiliated with party D). Redistricting then consists of choosing the changes in the partisan alignment of voters for each district, Δp_i, under the constraint that voters' partisan

[7] This can be considered a noisy estimate based on a single signal from the most recent election, or it could be constructed in the manner of *normal Democratic vote* by using information from several recent elections.

affiliations cannot be altered, thus $\sum_i \Delta p_i = 0$. It is worth emphasizing that our focus on the change, Δp_i, is deliberate in recognition that our bargaining process produces path dependence, as each new map is not laid down on a *tabula rasa* but made in reference to the previous lines.

If $\theta_i = p_i + \Delta p_i$ is the proportion of Republicans in district i on voting day, then the median voter's ideal point, $x_i(\theta_i)$, is that which solves $\theta_i I_R(x_i) + (1 - \theta_i) I_D(x_i) = \frac{1}{2}$. For each θ, there is a unique median, the median is strictly increasing in θ_i and $x_i\left(\frac{1}{2}\right) = 0$. Because the Republican Party wins if $d_i < x_i(\theta_i)$, the probability that a Republican wins the district is thus $f(\theta_i) = L(x_i(\theta_i))$. This is what Gul and Pesendorfer (2010) call the *District Outcome Function*. Their Lemma 1 establishes that as the leading party's support increases, its probability of winning increases, but at a decreasing rate: $f\left(\frac{1}{2}\right) = 0$ and $f' > 0, f'' < 0$ while $f(\theta_i) > \frac{1}{2}$.

Mapmaking consists of partitioning the set of precincts into N districts.[8] We presume that each member of the majority party has both office-holding and policy motives[9] such that the utility of a member of the majority party from district i is

$$U_i = f(p_i + \Delta p_i) + \gamma M\left(\frac{1}{N}\sum_{i\in N} f(p_i + \Delta p_i)\right). \tag{1}$$

The first term is the member's own probability of victory given the allocation of voters. This represents the office-holding motive. The second term represents the policy motive, and the parameter γ is the relative weighting. The second term centers on a function M, the argument of which is the expected seat share of the party drawing the map. [10] The function M represents the ability of the party to translate larger majorities into preferred policies. We assume $M' > 0$, implying that larger seat share translates monotonically into more preferred policies, but that $M'' < 0$, implying that the marginal value of an extra seat is declining in the size of the majority.[11]

We consider the case of cooperative bargaining among the members of the majority party and presume the Nash bargaining solution, in which case the map produced is that which satisfies

$$\begin{aligned} (2) \quad & max_{\{\Delta p_i\}_{i=1}^{N}} \Pi_{i\in N_R}(U_i - d_i) \\ & = max_{\{\Delta p_i\}_{i=1}^{N}} \Pi_{i\in N_R}\left(f(p_i + \Delta p_i) + \gamma M\left(\frac{1}{N}\sum_{i\in N} f(p_i + \Delta p_i)\right) - f(p_i)\right) \end{aligned}$$

[8] The role of precinct as the fundamental building block with which mapmakers work can be seen clearly in "The League of Dangerous Mapmakers" by Robert Draper (*The Atlantic*, October 2012), in which a well-known map consultant, in his presentation to legislators and staffers, "warns legislators to resists the urge to overindulge, to snatch up every desirable *precinct* within reach" [emphasis added].

[9] Fenno (1973) ascribes to legislators three motives: getting reelected, achieving influence within Congress, and making "good public policy." We include the first of these explicitly in our utility function, and the last is a function of securing a working majority for one's party and thus also included here. The second is beyond the scope of this model.

[10] We have also considered the case where the policy motive is based on maximizing the probability of achieving a majority. Propositions 1–3 continue to hold. Propositions 4 and 5 no longer hold perfectly, but simulations show that they hold almost everywhere.

[11] Microfounding these assumptions would necessitate adding within-party heterogeneity in policy preferences and an explicit policy process, which are beyond the scope of this paper.

AQ2 subject to

$$\sum_i \Delta p_i \;=\; 0,$$

where d_i is the disagreement point for member i. Among potential axiomatic bargaining solutions, we prefer the Nash approach because we favor scale invariance of utility, thereby ruling out other popular choices such as egalitarian or utilitarian solutions.

We choose to define the threat point as the current distribution of voters. Different states have different backup plans should the parties in charge fail to produce a map in time (Cox and Katz 2002); some remand the issue to commissions, some to courts. Frequently, the delays necessitate using the previous map for the first election past the due date or remanding to nonpartisan actors with little time for making big changes and little interest in making changes with systematic partisan effects, hence our selection of the previous map as the result of breakdown.

The first-order conditions imply the following:

$$\text{(3)} \qquad \left[\Psi \times \frac{\gamma}{N} M'\left(\frac{1}{N}\sum_{i\in N} f(p_i + \Delta p_i)\right)\right] f'(p_j + \Delta p_j) \;=\; \lambda, \forall j \;\in\; N_D$$

$$\text{(4)} \qquad \left[\Psi \times \frac{\gamma}{N} M'\left(\frac{1}{N}\sum_{i\in N} f(p_i + \Delta p_i)\right) + 1\right] f'(p_i + \Delta p_i) \;=\; \lambda, \forall j \;\in\; N_R$$

$$\text{(5)} \qquad \Psi \;=\; \sum_{l\in N_R} \Pi_{k\in N_R, k\neq l}\left(U_k - d_k\right).$$

The first set of conditions, obtaining for each of the seats currently held by the minority party, states that the increase in the probability that the majority party flips the district, f', times the value of an extra seat in passing policy, M', must equal the shadow price. The second set of conditions, obtaining for each of the seats currently held by the majority party, is similar but contains an extra $+1$ term inside the brackets indicating that the seat is valued not only for its effect on the majority but also directly by the member who would enjoy holding office.

From these, we can derive the following propositions:

PROPOSITION 1: *If* $p_i \;=\; p_j, i \;\in\; N_R, j \;\in\; N_D$, *then* $\Delta p_i \;>\; \Delta p_j$. *This follows directly from the first-order conditions and the assumption on* f''. *It implies that two districts, each won by a razor thin margin, one held by the majority party and one held by the opposition, will be treated differently, with more help sent to defending the marginal district already held than sent to flip the opposition district. Intuitively, this is because the representative of the majority-party district is represented in the bargain, while the potential challenger in the opposition district is not. Thus, the office-holding utility of the potential challenger is ignored in the collective bargain.*

PROPOSITION 2: *If* $p_i \;=\; p_j, i \;\in\; N_R, j \;\in\; N_D$, *then* $d\left(\Delta p_i - \Delta p_j\right)/d n_R \;>\; 0$. *In other words, the inequality in Proposition* 1 *will be stronger the larger the majority. This follows from the assumption about* M''. *Essentially, there is a tension between the office-holding motives and the policy-holding motives. The latter induces efficient* *direction of majority-party-aligned voters to where they add the greatest number of expected seats, while the former directs them to the districts currently held by the majority party. A larger majority reduces the marginal value of an additional seat, thereby tipping the balance toward shoring up existing seats.*

PROPOSITION 3: *On the contrary, the resulting inequality in Proposition 1 will be weaker in chambers and parties that place greater weight on policy motives relative to office-holding motives* $(larger\ \gamma)$. *This result follows directly from the first-order condition.*

AQ3 PROPOSITION 4: *If* $p_i < p_j, i, j \in N_R$, *then* $\Delta p_i > \Delta p_j$.

PROPOSITION 5: *If* $p_i > p_j, i, j \in N_L$, *then* $\Delta p_i > \Delta p_j$.

These last two results suggest that districts that are more competitive will receive more aid. This reproduces the intuition developed by Winburn (2008). Again, this follows from the assumption about f''.

B. *Simulation*

To visualize Propositions 1–5, we have numerically simulated our Nash bargaining solution. By plotting Δp_i against p_i, we are able to visually depict the discontinuity predicted by our model. Conveniently, the case of the unitary actor is nested in our model, allowing us to illustrate cleanly that the discontinuity from Proposition 1, which forms the heart of our empirical work, derives directly from relaxing that assumption. As the relative weight on policy increases, members with significant majority seats are willing to give up some of their supporters to assist the party in yielding a larger overall majority. In the limiting case, $\gamma \rightarrow \infty$, incentives are perfectly aligned, which mimics a unitary mapper, and our solution converges on the traditional, unitary actor, crack-and-pack results of Gul–Pesendorfer (2010). Analytically, this can be seen by noting that as $\gamma \rightarrow \infty$, the two first-order conditions become identical.

In the other limit, $\gamma \rightarrow 0$, sitting legislators are entirely self-serving. Utility gains occur only when friendly voters are reallocated from opposition districts by deeper packing and, thus, there is no motive to attempt to flip a district. Hence, there is a discontinuity in the manner in which competitive districts $(p_i \approx 0.5)$ are treated depending on whether they are currently part of the caucus $(p_i > 0)$ or not $(p_i > 0)$ (Figure 3). But when gamma is large, sitting legislators are willing to sacrifice their own chances of reelection to increase the expected party seat share. As a result, not only are competitive majority-party seats receiving support, but some of the opposition-held seats near the margin are targeted to be flipped and receive an influx of friendly voters (Figure 4). Thus, the discontinuity occurs at a very different place: between those seats that are being cracked—including some currently held by the opposition—and those that are being packed.

These simulations illustrate two other points from the model that are intuitive but not obvious. Notice that cracked districts are not perfectly equalized. Sitting

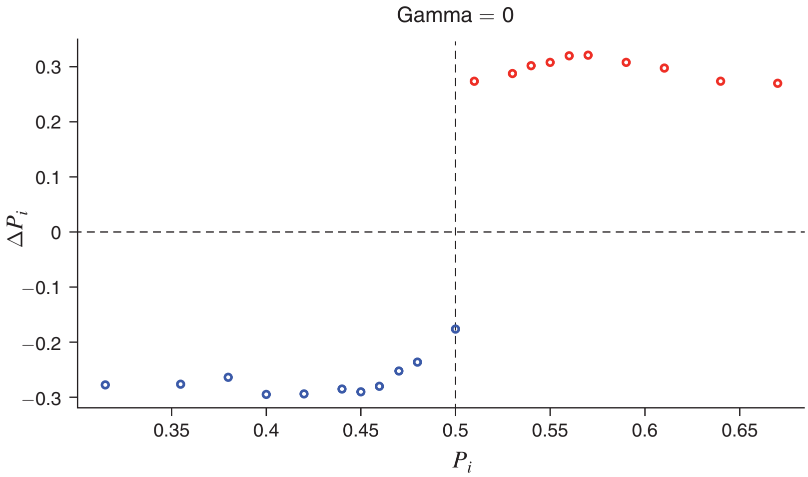

FIGURE 3. SIMULATED Δp_i-S FOR A 21-CHAMBER DISTRICT WHERE THE RED PARTY ENJOYS AN 11–10 MAJORITY AND CONTROL OF THE MAP

Note: When legislators care only about reelection, no districts are flipped and there is a huge discontinuity between marginally won and marginally lost districts.

legislators still defend their endowment of friendly voters. When $\gamma = 0$, every sitting legislator fights for an equal share of the friendly voters freed up by further packing (hence the horizontal slope of the right side of Figure 3). When $\gamma > 0$, individual legislators recognize that friendly voters contribute more to expected seat share if they are directed toward more competitive districts, and this leads to a downward slope. But even when gamma is sufficiently large to enable a caucus to attempt to expand, the cracked districts are not perfectly equalized: the slope of Δp_i versus p_i is flatter than -1; thus, $p_i + \Delta p_i$ is not equated across districts. This is another improvement on existing, history-free crack-and-pack results that predict equality across all cracked districts.

Second, notice that treatment of the districts held by the opposition, those to the left of $p_i = 0$, varies according to gamma. Every district that is set to be packed has the smallest geographically possible number for Δp_i as friendly voters are transferred out. But the districts targeted for flipping will have positive values. This is important to our empirical work. Critically, this breakpoint between crack and pack depends on gamma. Presumably, pooling chambers will result in a mixture of gammas. Thus, in some cases, there will be many targets on the left side of the line, and in other cases, relatively few. Moreover, because of varying geographic possibilities for redistricting and incumbent quality, those targets may not be neatly ordered by p_i, and thus, in practice, the left side of the graph is likely to be a jumble.

Finally, a word about geographic constraints. Sherstyuk (1998) proves that a requirement of contiguity alone does not constrain the ability of the mapper to partition voters into districts. However, common TDPs such as compactness, respect for communities of interest, respect for political subdivisions, and maintaining the core of a district mean that geography does likely play a role in constraining

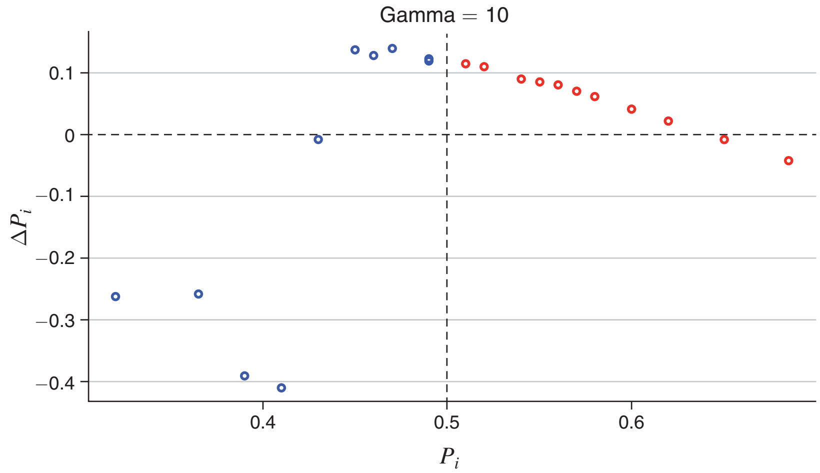

FIGURE 4. SIMULATED Δp_i-S FOR A 21-CHAMBER DISTRICT WHERE THE RED PARTY ENJOYS AN 11–10 MAJORITY AND CONTROL OF THE MAP

Notes: When legislators are group-minded, they do seek to flip opposition districts. In some cases, the safest incumbents will even be willing to outright lose voters. The boundary between those opposition districts which are targeted and those which are packed is situation dependent.

the extent to which a mapper can transfer friendly voters between districts. The ability to accurately identify the locations of copartisans may also matter. For instance, if voters cannot be located more accurately than their precinct and precincts are never more lopsided than 80–20, then a district cannot be packed beyond 80–20.

A simple modeling of geographic constraints would be to suppose that the majority-party legislators first identify which opposition districts to pack and then apportion the friendly voters that they can extract from these districts among the remaining districts according to the bargaining process above. This adjustment does allow for the possibility that the number of transportable voters depends on which districts are targeted for improvement, but we presume that reallocation of friendly voters among targeted districts is unconstrained. It is simple to demonstrate that our results hold in this case. Given the ability to transmit friendly voters from one district to a nonadjacent district by a chain of swaps and the general conclusion from scholars of gerrymandering that TDPs are not an important constraint, we feel that this is a reasonable approximation. A more detailed model would suppose that certain districts are more difficult to improve than others because of the distribution of voters. We have explored an approach where each district is characterized by a range of feasible values for Δp_i. This range is exogenously determined by the partisan distribution of nearby precincts and which swaps would not violate TDPs to the extent that a lawsuit would be found meritorious. For certain assumptions of the correlation between $Deltap_i$ and p_i, and with an extension of the Nash bargaining concept, we can show that the main results of the paper go through.

IV. Data Sources and Preparation

To test our propositions, we obtained historical voting records of precincts across the United States from the Stanford Election Atlas (Rodden and Ansolabehere 2011). The records indicate how individuals within precincts voted over the period of 2004 to 2008 for state gubernatorial, attorney general, secretary of state, controller, treasurer, insurance commissioner, congress, assembly, and senate elections. In addition to state-level elections, voting records for the presidential election in 2008 are also included. Normal Democratic vote share (*NDV*) is estimated for a precinct by averaging Democratic vote share across all of the aforementioned elections.[12] From Klarner's (2018) state legislative returns dataset, we know which party controls each district.

Each precinct's voter data is linked to a shapefile, a geospatial vector data format for GISes. Along with our precinct-voting-level GIS data, we extract shapefiles for the state legislative district lines in 2006 and 2015 from the US census TigerLines database for the lower and upper chambers of each state. The redrawing of district lines is typically conducted after each census to account for changes in population estimates. As a result, our 2006 state legislative district lines represent districts over the period of 2001 to 2010, and our 2015 state legislative district lines represent districts over the period of 2011 to 2020. Unfortunately, prior redistricting waves are not available as GIS shapefiles, so we are limited to these two waves. Fortunately, we nonetheless have a pair of maps from which to calculate changes in partisan vote shares of districts as the map is redrawn.

We combine our precinct-level *NDV* data with our upper and lower chamber state legislative district lines by first converting the geospatial projections of our data into a common coordinate reference system (CRS) through the Environmental Systems Research Institute's (ESRI's) ArcMap software. Once projected into a common CRS, we use ESRI's intersect tool to find the shared area between our precinct-level data and the respective upper and lower chamber state legislative district lines for 2006 and 2015. More precisely, we find the percentage of the area in square kilometers of each precinct that falls within a district, to assign precincts to districts.

District-level *NDV* is then computed as the population- and area-weighted *NDV* of each of the assigned precincts. For example, if precinct i has an *NDV* of 0.6 with 1,000 voters and is geographically split 50-50 between districts j and k, district j will receive 500 voters with an *NDV* of 0.6 and district k will receive 500 voters with an *NDV* of 0.6. As a result of precincts being our smallest measurement unit, we assume a homogeneous geographical distribution of voters within each precinct. District-level *NDV* is computed in this manner for each of the districts in our lower and upper chamber state legislative district lines for 2006 and 2015. To account for potential changes in district names across the redistricting wave in 2010, we compare how the lower (upper) district lines change by analyzing the population-weighted

[12] We have also used Cooperative Congressional Election Study data on demographics and vote choice to model the demographically predicted vote choice, which we call expected democratic vote (EDV). We then use census data to calculate EDV at the block group level. The results are similar to what we report for *NDV,* but weaker, as demographics are merely a portion of vote choice and mappers have the more complete data on vote choice.

area overlap between lower (upper) chamber lines in 2006 and 2015. We intersect the lower (upper) chamber 2006 and 2015 state legislative district lines to find the shared common area between each district in 2006 with each district in 2015. The common areas are then weighted by the share of the original population in each district to estimate the percentage of voters passed on from an original district i to each of the 2015 districts within the same state as district i. We can thus identify which "offspring" district k corresponds to which "parent" district i. The offspring is that which takes on the largest fraction of voters originally in the parent. We discuss our matching procedure in detail in the Appendix.

V. Empirical Work

Our primary sample consists of 25 of the 28 states in which a single party controlled redistricting during the 2010 wave (18 Republican, 7 Democrat). Generally speaking, control by one party requires that party be able to pass normal legislation without any votes from the other party. This usually requires a majority in both chambers and, where the map can be vetoed, the governorship as well.[13] We also run placebo tests using the chambers where both parties were required to approve the map or the process was handed to an independent commission. Each district is a data point. As described in Section IV and detailed in the Appendix, we have matched districts before and after the redistricting.

Using our previously described data on *NDV*, we calculate the change in *NDV* between the parent and offspring. We then adjust the sign so that in each case, a positive value means that the district is becoming more favorable for the party in control of redistricting. This ΔNDV_i, a measure of the extent to which the district is made more or less favorable to the party in control, serves as our dependent variable. For the chambers where control was either shared or delegated to an independent commission, we use the unadjusted change in *NDV*.

Our theory implies that districts currently under the control of the party conducting the redistricting will be treated differently than similarly competitive districts currently in opposition because the candidates likely to run for the majority party in the out-districts are not represented in the bargaining process. We look for evidence of this hypothesis using first a regression discontinuity design and then a cross-sectional regression that allows for both intercepts and slopes to vary by pre-redistricting control. In each case, we focus on those districts that are within shouting distance of contestation (electoral margin < 0.25). Both analyses exhibit similar results and corroborate our theory. In each analysis, the result disappears when we run the baseline model using chambers where no single party controls the map.

Figure 5 shows the plot of the binned data with a fitted second-order polynomial. The three most important results are all clear from this figure: the discontinuity showing extra support for districts already under control, the downward slope on the right-hand side indicating that help for sitting copartisans depends on the competitiveness of the

[13] Data on control of redistricting are from Justin Levitt's Web site *All About Redistricting*. Levitt explains the institutional variation among states. Alaska, Massachusetts, and Nebraska were dropped from our sample due to an inability to match districts across waves or lack of data on partisan identities of incumbents.

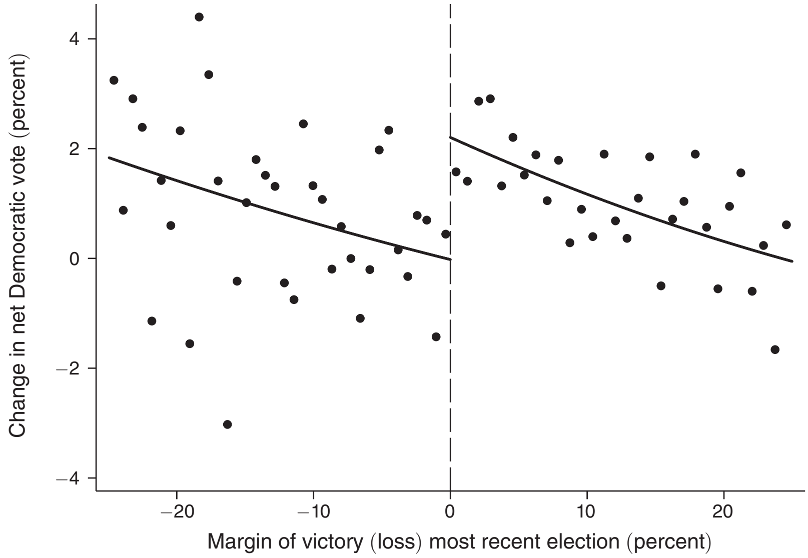

FIGURE 5. RD PLOT: PRESERVING TAKES PRECEDENCE OVER FLIPPING

Notes: The discontinuity between held and targeted seats is evident. Seats held by the majority (RHS) receive greater infusions of friendly voters the more competitive that the most recent election was. Note also that margin of victory explains partisan shifts better for districts already held by the majority.

district, and the much greater variance on the left-hand side, which is consistent with our prediction of a shifting boundary between those districts to be packed and those to be targeted for flipping. (See the simulations in Section IIIB.)

The regression discontinuity is estimated using Calonico, Cattaneo, and Titiunik's (2014) *rdrobust* command in Stata, with defaults for bandwidth and local polynomial order (1). We show in Table 1 that the results are robust to these decisions and to the sample restriction. The results suggest that, depending on the specification, a marginal district will get between 2.0 and 2.7 percentage points more help if it is represented by a sitting legislator. The final column switches the sample from those maps controlled by a single party to maps jointly influenced by both parties (e.g., because the legislature and the governor are of different parties and the majority is not veto-proof) or delegated to an independent commission, such as in California. As expected, the central result is not present unless the map is controlled by a single party.

Column 1 of Table 2 shows results for the corresponding OLS regression. The constant term shows that the marginal opposition-held seat gets zero help on average. Meanwhile, the marginal seat held by the party drawing the lines gets 2.2 percentage points of help on average, a magnitude closely in line with the regression discontinuity results. We also see the expected relationship whereby the extent of assistance to own-party incumbents declines as the seat becomes safer. The estimated coefficients suggest that seats with margins of victory above 24 percent receive no further assistance. The surprise from Figure 5 is confirmed, as we see no evidence that action toward opposition held seats is systematically related to the

Table 1—Regression Discontinuity Results

Control of map	Single party (1)	Single party (2)	Single party (3)	Single party (4)	Single party (5)	Bipartisan + commissions (6)
RD_estimate	0.0247	0.0241	0.0266	0.0213	0.0200	−0.001
Observations	1,925	1,674	1,925	1,925	1,925	781
Eff observations	512	501	982	927	1,314	243
Kernel type	Triangular	Triangular	Triangular	Triangular	Triangular	Triangular
BW type	mserd	mserd	mserd	Manual	Manual	Manual
Conventional std. error	0.009	0.009	0.010	0.007	0.006	0.009
Conventional p-value	0.007	0.009	0.006	0.002	0.000	0.912
Robust p-value	0.015	0.024	0.011	0.007	0.005	0.847
Order loc. poly. (p)	1	1	2	1	1	1
Order bias (q)	2	2	3	2	2	2
BW loc. poly. (h)	0.055	0.053	0.106	0.100	0.150	0.063
BW bias (b)	0.090	0.084	0.154	0.100	0.150	0.111
Abs(electoral margin) ≤	0.25	0.20	0.25	0.25	0.25	0.25

Notes: This table reports the results of five separate specifications over which three parameters are varied: the order of the local polynomial (p), the bandwidth of the local polynomial (h), and the extent to which the sample is limited to competitive districts (absolute value of electoral margin). All of these specifications are significant at conventional levels, and these estimation choices have relatively little effect on the magnitude of the effect. In the first five columns, a single party controls the map. The sixth column switches the sample to cases where control is shared between the parties or delegated to an independent commission. The coefficient is practically and statistically indistinguishable from zero, showing that competitive districts are not treated differentially according to which party won if there is no asymmetry in control.

margin of recent defeat. The R^2 is likewise an astonishingly low 1 to 3 percent, on which we will comment more later.

We then split the sample into chambers that are competitive and those that are not, defining competitive as neither party has a two-thirds supermajority. In our sample of chambers whose redistricting is under unified partisan control, 13 out of 25 upper chambers and 15 out of 25 lower chambers are competitive by this definition, giving us a roughly equal split of districts and sample size across these subsamples. We find that the result in question—the discontinuity according to whether the seat is currently held—is twice as large in the competitive chambers as it is in the uncompetitive chambers. Again, this is precisely as predicted by our bargaining theory (Proposition 2). In uncompetitive chambers, the value of the public good is lesser (large majorities ensure that favorable policy is not dependent on adding an additional seat); thus, there is greater emphasis on individual incentives.

Noting that chamber size, political experience, and term of office might affect the bargaining process, our second split is between lower and upper chambers. Lower chambers are, on average, three times larger than their upper chamber counterparts. They are also far more likely to have 2-year terms instead of 4-year terms: 44 of the 49 lower chambers have 2-year terms, whereas 38 of the 50 upper chambers have 4-year terms. Nonetheless, we find no significant difference between lower and upper chambers in the magnitude of the effect. Finally, we have run both the regression discontinuity and the OLS for the sample of maps drawn by independent and bipartisan commissions (Table 1, column 6; Table 2, column 4). In each case, the result in question disappears, consistent with the theory that this arises from intraparty bargaining.

Table 2—OLS Results

Dependent variable	ΔNDV	ΔNDV	ΔNDV	ΔNDV	ΔNDV	ΔNDV
Chambers	All	Uncompetitive	Competitive	Competitive	Upper	Lower
Control of map	Single party	Single party	Single party	Bipartisan + commissions	Single party	Single party
	(1)	(2)	(3)	(4)	(5)	(6)
Own seat	0.0219	0.0128	0.0288		0.0256	0.0214
	(0.00451)	(0.00602)	(0.00654)		(0.00889)	(0.00525)
Margin of victory	0.0739	0.0241	0.0990	0.00703	0.118	0.0587
	(0.0288)	(0.0412)	(0.0394)	(0.0612)	(0.0581)	(0.0331)
Own seat × margin of victory	−0.165	−0.0531	−0.243		−0.255	−0.136
	(0.0357)	(0.0485)	(0.0509)		(0.0704)	(0.0414)
Held by Democrat (Dem)				0.000300		
				(0.0115)		
Dem × margin of victory				−0.0297		
				(0.0985)		
Constant	−0.000677	0.00206	−0.00199	−0.0184	−0.00527	0.000756
	(0.00349)	(0.00498)	(0.00477)	(0.00824)	(0.00726)	(0.00396)
Observations	1,948	862	1,086	663	568	1,380
R^2	0.014	0.007	0.024	0.000	0.028	0.013

Notes: Simple OLS regressions are in line with the RD results reported in Table 1 and highlight three additional results. First, the slopes with respect to electoral margin are significant and of opposite sign, indicating, as expected, that friendly voters are transferred from uncompetitive districts to competitive districts. Second, the result is confined to competitive chambers. Third, there is no difference between upper and lower chambers despite the systematic differences in term length and chamber size that might have affected bargaining. As with the RD, the effect is not present when no single party controls the map. Standard errors in parentheses.

The miniscule values of R^2, never topping 3 percent in any specification, suggest that the shift in the partisan composition of seats is largely unexplained by the sort of maximization that dominates the classic crack-and-pack theories. We have argued that this is consistent with a bargaining model in which a large set of incumbent politicians have the power to threaten reversion to the status quo, thereby ensuring that sacrifices for the common good of increased seat share are rare. As a result, the current map is somewhat sticky, and maps display path dependence that results in deviation from the unitary actor optimum. One alternate explanation is that legal constraints constitute an important source of friction inhibiting the optimal map, especially the traditional redistricting principles highlighted by Winburn (2008).

To address this possibility, we calculate measures of population-weighted overlap to determine the fraction of its original voters that a district retained during the redistricting process. We find that districts held by the party in control of redistricting retain a significantly larger fraction of seats (Table 2, first row of coefficients). The average overlap between parent and offspring is 73.4 percent in our sample. This, along with the low R^2 in Table 2, is strong evidence that whatever constrains shifts in partisan alignment is not a legal constraint on moving district boundaries.

AQ4

VI. Discussion

We have argued that existing theories of partisan gerrymandering are likely to overestimate the degree to which the mapmaker can pursue seat-share maximization when drawing the new map. The fixed set of friendly voters is a scarce resource over which members of the majority party have only partially aligned preferences. These voters simultaneously provide the public good of expected seat share and the rivalrous private good of one's own chance of reelection. Thus, we propose to replace the unitary decision-maker with a bargaining framework. This shift delivers two important impediments to the maximization of expected seat share. The first is the emphasis on those districts whose representatives are an active part of the bargaining process. The second is the sense that current representatives have some form of property rights over their current districts and must agree to trade them away. As a result, the existing map becomes an important point of departure. This introduces a role for history and inertia in the pursuit of seat-share maximization.

But our presumed threat point is a strong requirement: the Nash result that every individual legislator must receive surplus beyond the disagreement point seemingly suggests an underlying structure in which any individual legislator can veto a map. Why, when moving away from the concentration of power in an individual, should we move to the other end of the spectrum? It is likely that the actual process of drawing the lines is done by a small committee in extensive discussion with the broader membership, representing at least some concentration of power. Is it not out of the question that some members should be called upon to sacrifice? Does not the party have the fungible resources and longevity to enable trade credit for such sacrifices?

Perhaps one ought to view our model as representing the other end of a spectrum and thus usefully illustrating the effects of moving some distance in that direction. And yet, our empirical results suggest that only a tiny fraction of the shifting of voters produces a net change in the partisan balance of districts in the pattern that would imply improvements in expected seat share. This might be evidence of the difficulties in effecting net changes in vote share that result from the fact that most precincts are mixed and contiguity prevents grabbing distant voters who might offer the needed concentration. We believe it is also evidence of the inertia that comes from decentralized bargaining with broadly distributed power to revert to the status quo.

One puzzle that remains is why, among opposition districts, the change in the partisan vote share is not a clear function of the competitiveness of the district. One possibility for this confusion has already been discussed: the pooling of chambers pools differing cutoffs. Another possibility for the asymmetry is that the bargaining process requires defending all majority-party districts but that opposition districts are picked off solely based on opportunity. The lack of a relationship could then be the product of either of two branches. Either opportunities are not meaningfully correlated with the competitiveness of the district, or the flipping of opposition districts is not pursued by the transference of friendly voters, possibly because that requires that those voters be donated from a district already held.

Opportunity might arrive in the form of retirements, scandals, and the possibility of pairing opposition incumbents, none of which would appear in our data.

Scandals are likely equally distributed across districts and would thus simply be noise, obscuring any existing relationship but not fundamentally explaining the lack thereof. Opportunities to pair incumbents would seem to be more useful if the district is actually winnable demographically; there is no partisan gain in pairing two incumbents deep in opposition territory. Retirements may be more likely in close districts. All of this contributes to the greater variance in the treatment of opposition districts that is evident in the regression discontinuity plot.

In sum, the net shifts in the partisan composition of voters in state legislative chambers in the 2010 wave are remarkably muted when compared to the predictions of models based on unitary mappers. We suggest that this could be explained as the result of a bargaining process in which sitting legislators of the majority party enjoy broadly dispersed power to default to the existing map. Our decentralized bargaining theory further predicts a discontinuity in the treatment of competitive districts already held by the majority party and those currently held by the opposition. Moreover, this discontinuity ought to be greater in chambers where the majority is large. We find support for both of these hypotheses. We thus have specific empirical support for our theoretical proposition that redistricting is a bargaining process that privileges current members of the chamber.

Appendix

In the text, we have described the process of matching *parent* and *offspring* districts from successive waves. Ideally, this mapping would be one-to-one and onto. Unfortunately, there is no single, obvious method by which to produce a mapping that is one-to-one and onto, and yet some choices in this mapping method affect the outcome. Nonetheless, we believe that we have the proper mapping and that our results are robust to alternative appropriate mappings and thus relegate this more detailed explanation to the Appendix.

Our mapping procedure was this: for every parent district from the prior wave, assign as its offspring district that district from the successor wave to which it (the parent) donated the largest number of voters. Thus, if district A were split across districts A', B', C' 20-45-35, B' would be designated as the offspring of A. To calculate the change in *NDV*, we would subtract the *NDV* of A from that of B'. For the regressions, we would be pairing this ΔNDV with the most recent electoral margin in parent district A.

One can imagine several other mapping procedures. The simplest change would be to match parents to offspring rather than the other way around. That is, for every offspring, assign to it as parent that district from which the greatest fraction of the offspring is derived. In the overwhelming majority of cases, this change makes no difference. Consider Table A1, illustrating a hypothetical set of three districts.[14] In this case, the (*parent*, *offspring*) pairs are (A,B'), (B,A'), (C, C') no matter which

[14] Notice that while the rows must sum to 100 percent, the columns do not, as a result of differential population growth rates. For example, if district X grows much more slowly than the rest of the state, then 100 percent of district X would be insufficient to furnish the full population of successor district X', which would need some fraction, say 10 percent, of district Y, in which case the X' column would sum to 110 percent.

Table A1—Three-District Example 1

Percent of parent in offspring		Offspring A’	B’	C’
Parent	A	20	45	35
	B	75		25
	C		52	48

Table A2—Three-District Example 2

Percent of parent in offspring		Offspring A’	B’	C’
Parent	A	20	45	35
	B	75		25
	C		48	52

Table A3—Five-District Example

Percent of parent in offspring		Offspring A’	B’	C’	D’	E’
Parent	A	20	45		35	
	B	60		40		
	C		70	30		
	D			20	80	
	E			5		95

direction is chosen for the matching. But if we consider the slightly modified example in Table A2, we now see that the direction of the matching matters.

Matching offspring to parents produces (A,B’), (B, A’), (C, B’), whereas matching parents to offspring produces (A’, B), (B’, C), (C’, C). This also shows how the matching is neither one-to-one nor onto. In the second case, offspring B’ happens to be the largest recipient from both A and C. Likewise, parent C is the largest donor to both B’ and C’. There are essentially three ways of dealing with this. The first is to accept the match as is. The second is to remove multiple matches according to some priority and rematch the leftover parents and offspring according to some alternate rule. The third is to remove the multiple matches without rematching. None are ideal.

The first method results in a partially complete map in that either all the parents or all the offspring are used, but not both. The strength of this approach is that a clear and consistent relationship between parent and offspring is maintained. The third method similarly maintains a clear relationship between the parent and offspring of the maintained matches, with the added benefit of avoiding double-use of any parents or offspring, but at the cost of an incomplete map and a choice over how to prioritize among multiple matches.

Achieving a one-to-one and onto mapping requires the second method, which necessitates another method of matching. [15] Unfortunately, in most cases, one is left matching parents and offspring that have zero overlap. To understand why, consider the five-district example in Table A3. The assignment of offspring to parents results in both A and C wishing to claim B'. As C clearly has greater claim to B' than A does, having donated 70 percent rather than 45 percent, we assign B' to C and search for a new match for A. Unfortunately, both A' and D' have already been assigned to B and D, respectively. The unassigned offspring is C', with which A shares no overlap. In this particular instance, one might argue that if we were to assign the contested offspring B' to A, then C' could be assigned to C, thus ensuring that the secondary pairing also enjoys nonzero overlap. We have experimented with such schemes and found that they solve relatively few cases and at the cost of significantly reducing the overlap of the first match.

Our model considers the effect of incumbent preferences on the district in which they run. As they must be a resident to run in the district, a new district with no overlap is not an incumbent-relevant offspring. As such, we wouldn't expect our theory to be relevant to such matches. Indeed, our core results go through with either offspring-to-parent or parent-to-offspring matching and either allowing many-to-one matches or keeping only the strongest such match. However, when we attempt to rematch the remaining parents and offspring, the resulting noise overwhelms the result. Thus, despite the desirability of a one-to-one and onto mapping, the lack of a clear relationship between the rematched parents and offspring makes it clear that these districts ought to be left out of the analysis.

AQ5

REFERENCES

Acemoglu, Daron, Simon Johnson, and James A. Robinson. 2005. "Institutions as a Fundamental Cause of Long-Run Growth." In *Handbook of Economic Growth*, Vol. 1A. 385–472. Amsterdam: Elsevier.

Aghion, Philippe, Alberto Alesina, and Francesco Trebbi. 2004. "Endogenous Political Institutions." *Quarterly Journal of Economics* 119 (2): 565–611.

Berry, Christopher R., Barry C. Burden, and William G. Howell. 2010. "The President and the Distribution of Federal Spending." *American Political Science Review* 104 (4): 783–99.

Besley, Timothy, and Ian Preston. 2007. "Electoral Bias and Policy Choice: Theory and Evidence." *Quarterly Journal of Economics* 122 (4): 1473–1510.

Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik. 2014. "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs." *Econometrica* 82 (6): 2295–2326.

AQ6 **Carson, Jamie L., Michael H. Crespin, Charles J. Finocchiaro, and David W. Rohde.** 2007. "Redistricting and Party Polarization in the U.S. House of Representatives." *American Politics Research* 35 (6): 878–904.

Chen, Jowei, and Jonathan Rodden. 2013. "Unintentional Gerrymandering: Political Geography and Electoral Bias in Legislatures." *Quarterly Journal of Political Science* 8: 239–69.

Cho, Wendy K. Tam, and Yan Y. Liu. 2016. "Toward a Talismanic Redistricting Tool: A Computational Method for Identifying Extreme Redistricting Plans." *Election Law Journal* 15 (4): 351–66.

Coate, Stephen, and Brian Knight. 2007. "Socially Optimal Districting: A Theoretical and Empirical Exploration." *Quarterly Journal of Economics* 122 (4): 1409–71.

Cox, Gary W., and Jonathan N. Katz. 2002. *Elbridge Gerry's Salamander: The Electoral Consequences of the Reapportionment Revolution*. Cambridge, U.K.: Cambridge University Press.

[15] Unlike Congressional redistricting, we almost never have to deal with the loss or gain of a seat during the 2010 wave. The exception is the NY Senate, which added a sixty-third seat. In that case, we chose to allow one of the offspring to remain unmatched.

Daley, David. 2016. *Ratf**ked: Why Your Vote Doesn't Count*. New York: W. W. Norton & Company.

Fryer, Roland G., Jr., and Richard Holden. 2011. "Measuring the Compactness of Political Districting Plans." *Journal of Law and Economics* 54 (3): 493–535.

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

Gul, Faruk, and Wolfgang Pesendorfer. 2010. "Strategic Redistricting." *American Economic Review* 100 (4): 1616–41.

Owen, Guillermo, and Bernard Grofman. 1988. "Optimal Partisan Gerrymandering." *Political Geography Quarterly* 7 (1): 5–22.

McGhee, Eric. 2014. "Measuring Partisan Bias in Single-Member District Electoral Systems." *Legislative Studies Quarterly* 39 (1): 55–85.

Rodden, Jonathan. 2019. *Why Cities Lose: The Deep Roots of the Urban-Rural Political Divide*. New York: Basic Books.

AQ7

Rodden, Jonathan, and Stephen Ansolabehere. 2011. "Stanford Election Atlas." https://projects.iq.harvard.edu/eda/data.

Sabouni, Hisam, and Cameron Shelton. 2020. "State Legislative Redistricting: The Effectiveness of Traditional Districting Principles in the 2010 Wave." *Election Law Journal* 20 (2): 198–214.

AQ8

Sabouni, Hisam, and Cameron A. Shelton. 2022. "Replication data for: Gerrymandering in State Legislatures: Frictions from Axiomatic Bargaining." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.38886/E142061V1.

Sherstyuk, Katerina. 1998. "How to Gerrymander: A Formal Analysis." *Public Choice* 95 (1/2): 27–49.

Shotts, Kenneth W. 2001. "The Effect of Majority-Minority Mandates on Partisan Gerrymandering." *American Journal of Political Science* 45 (1): 120–35.

Winburn, Jonathan. 2008. *The Realities of Redistricting: Following the Rules and Limiting Gerrymandering in State Legislative Redistricting*. Plymouth, U.K.: Lexington Books.