Gerrymandering in State Legislatures: Frictions from Axiomatic Bargaining

Preliminary, please do not quote or cite without permission

Hisam Sabouni
Claremont Graduate University

Cameron Shelton
Claremont McKenna College

May 2018

Abstract

The most prominent existing theories of partisan redistricting postulate unitary actors maximizing expected seat share. Yet the partition of a fixed supply of friendly voters necessarily implies a tragedy of the commons as each individual legislator would prefer that improvements in the expected majority come from a colleague's district rather than endangering his own chances of reelection. We recast partisan redistricting as a bargaining game among the sitting representatives of the party that controls the chamber. The status quo is the threat point, explaining why changes are frequently minor, even after shifts of partisan control. This bargaining framework implies that highly competitive districts will receive more help from the redistricting if they are already represented by the party in charge. Using precinct level data regarding partisan affiliation, we test this prediction using a regression discontinuity design and find support.

I. Introduction

While the study of the boundaries of US Congressional districts has a long history, public interest has recently turned to US Statehouses. This is, in part, because a gridlocked US Congress has shifted legislative activity to the states and, in part, because it has been recognized that control of state legislatures enables a party to draw congressional district boundaries.[1] Recent attention to this longstanding truth seems to have been the result of the extreme success of a deliberate and well-funded Republican strategy to capture state legislatures in 2010 so as to control the subsequent decennial redistricting in 2011 (Daley 2016). Not only did Republican majorities in state legislatures draw the district lines for their Congressional delegations, they also drew district lines for their own state legislatures, potentially entrenching this source of political power, to the chagrin of Democrats who have mounted legal challenges.

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[2] 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.

There are two prominent such 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 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 map-maker will then *slice* the remaining voters in the same way, *covering* the most ardent foes with a slightly larger set of

[1] It also multiplies the number of maps under study from 1 to 99, which is a boon to statistical analysis.

[2] 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.

the most ardent supporters. And so on until the final district is a single 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 the median is a friendly voter. The size of this "overbite" depends on the degree of volatility in voter preferences. Friedman and Holden 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.

Yet parties are not unitary. Members face a complex combination of individual and collective goals. The partitioning of a fixed electorate with limited co-partisans implies a dilemma for any individual member. Optimizing 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 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 optimizes seat share is unlikely to emerge.

While most states pass the new map through standard legislative procedure, the party in power typically introduces a map that has already been agreed upon amongst members. Written accounts of the intra-party 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.[3] 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 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

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

on a map before the deadline, the state might use either a backup commission or remand the issue to the state court. In some cases, these backstops will share the partisan orientation of the legislature itself, but in most cases, the alternative is less partisan. To capture this lesser partisanship and the sense that caretaker commissions with little time and expertise are less likely to radically redraw the map, 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 an important reference point. We believe this is important to explain what seems to be a high degree of inertia in the lines.

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.

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 towards 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 overturns 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 electorate each appears 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.[4] But

[4] 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 they are the 50-50 seats, they will so remain), this means

because the latter is represented in the bargain, that district's changes will be more favorable. Our model thus delivers multiple testable implications for the relationship between a district's margin of victory in the most recent election and the change in expected vote share delivered when the lines are redrawn.

Using precinct-level data, we calculate the change in normal democratic vote that results for each district from the 2010 redistricting wave. This requires 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 map-makers 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, friendly voters are transferred on net from safe seats to competitive seats. Third, despite these statistically significant results in the expected direction, the R-squared is shockingly low. That suggests the great majority of the reallocation of voters is neutral from a partisan perspective. This supports the idea that whatever inertia exists is particular to the partisan rebalancing of districts. Given that, prior to 2017, the courts routinely refused to comment on partisan gerrymanders, it is unlikely this reluctance to shift the partisan balance comes from the threat of legal action. We suggest it is consistent with frictions in the bargaining process.

The remainder of our paper consists of four sections. First, we develop a simple bargaining model. Second, we present our data. Third, we conduct our empirical work. Finally, we conclude with discussion of the implications and ideas for further work.

---

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.

II. Theory

In this section, we 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 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 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(·). 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, 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.[5] Prior to the current round of redistricting, each district *i* is endowed with a set of voters characterized by a fraction p_i

[5] 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.

affiliated with party *R*. Redistricting then consists of choosing the changes in the partisan alignment of voters for each district, Δp_i, under the constraint that voter's partisan affiliations cannot be altered, thus $\sum_i \Delta p_i = 0$.

If $\theta = p_i + \Delta p_i$ is the proportion of Republicans in the district on voting day, then the median voter's ideal point, x(θ), is that which solves: $\theta I_R(x) + (1-\theta) I_D(x) = 1/2$. For each θ, there is a unique median, the median is strictly increasing in θ, and x(½) = 0. Because the Republican party wins if d<x(θ), the probability a Republican wins the district is thus: $f(\theta) = L(x(\theta))$. This is what GP 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(½)=0 and f'>0, f''<0 while f(θ) > ½.

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

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

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 factor 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. 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. [6]

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

$$\max_{\{\Delta p_i\}_{i=1}^N} \prod_{i\in N_R} (U_i - d_i) = \max_{\{\Delta p_i\}_{i=1}^N} \prod_{i\in N_R} \left(f(p_i + \Delta p_i) + \gamma M\left(\frac{1}{N}\sum_{i\in N} f(p_i + \Delta p_i)\right) - f(p_i) \right)$$

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

$$s.t. \sum_i \Delta p_i = 0$$

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 points 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, hence our selection of the previous map as the result of breakdown.

The first order conditions imply the following:

$$\left[\Psi * \frac{\gamma}{N} M'\left(\frac{1}{N}\sum_{i\in N} f(p_i + \Delta p_i)\right) + 1\right] f'(p_i + \Delta p_i) = \lambda, \quad \forall \mathrm{i} \in N_R$$

$$\left[\Psi * \frac{\gamma}{N} M'\left(\frac{1}{N}\sum_{i\in N} f(p_i + \Delta p_i)\right)\right] f'(p_j + \Delta p_j) = \lambda, \quad \forall \mathrm{j} \in N_D$$

$$\Psi = \sum_{l\in N_R} \prod_{k\in N_R, k\neq l} (u_k - d_k)$$

From which 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: This resulting 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 γ). This result follows directly from the first order condition.

Proposition 4: $p_i < p_j, i, j \in N_A$ then $\Delta p_i > \Delta p_j$

Proposition 5: $p_i > p_j, i, j \in N_B$ then $\Delta p_i > \Delta p_j$

These last two results suggest that districts which are more competitive will receive more aid. This reproduces the intuition developed by Winburn (2007). Again, this follows from the assumption about f''.

III. Data Sources and Preparation

To test our propositions, we compile historical voting records of precincts across the United States from the Stanford Election Atlas (Rodden and Ansolabehere). 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.

Each precinct's voter data is linked to a shapefile, that is a geospatial vector data format for geographic information systems (GIS). Along with our precinct voting level GIS data, we extract shapefiles for the state legislative district lines in 2006 and 2015 from the U.S. Census TigerLines database for the lower and upper chambers of each state (https://www.census.gov/geo/maps-data/data/tiger-data.html). The re-drawing 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.

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 Environmental Systems Research Institute's (ESRI) 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 must 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) chamber district lines change by analyzing the population weighted 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*. District *i* is matched from 2006 to 2015 by following the district *k* that takes on the largest fraction of voters originally in district *i.*

IV. Empirical Work

Our sample is the 27 states in which a single party controlled redistricting during the 2010 wave (20 Republican, 7 Democrat). Each district is a data point. As described in section III, we have matched districts before and after the redistricting.

Using our previously described data on normal democratic vote (NDV), we calculate the change in normal democratic vote between the parent and offspring. We then adjust the sign so that in each case, a positive value means 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.

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 cross-sectional regression that allows for both intercepts and slopes to vary by pre-redistricting control and then a regression discontinuity design. Both analyses exhibit similar results. In each case, we focus on those districts that are within shouting distance of contestation (electoral margin < 0.2) and are not completely eviscerated ($|\Delta NDV| < 0.2$).

The regression discontinuity is estimated using Calonico et al's *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 1.4 and 2.9 percentage points more help if it is represented by a sitting legislator. Figure 1 shows the plot of the binned data with a fitted 4^{th} 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 co-partisans depends on competitiveness of the district, and the lack of slope on the left-hand side suggesting, counterintuitively, that targeting of opposition districts is not systematically related to competitiveness or not achieved by shifting the partisan composition of voters.

Column 1 of Table 2 shows results for the OLS regression counterpart. 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 1.5 percentage points of help on average, a magnitude broadly 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 16% receive no further assistance. The surprise from figure 1 is confirmed as we see no evidence that action towards opposition held seats is systematically related to the margin of recent defeat. The R-squared is likewise an astonishingly low 0.3%, 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 2/3 super-majority. In our sample of chambers whose redistricting is under unified partisan control, 13/27 upper chambers and 15/26 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 present only in the subsample of uncompetitive chambers. Recall 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 favorable policy is not dependent on adding an additional seat) thus there is greater emphasis on individual incentives. Nonetheless, the portion of variation explained remains miniscule.

Noting that chamber size, professionalism, and term of office might affect the bargaining process, our second split is between lower and upper chambers. We find the effect is concentrated in the lower chamber rather than the upper chamber. 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. This high degree of collinearity makes it difficult to disentangle whether this form of gerrymandering occurs in lower chambers because they are small or because they have short terms. (The answer could also, of course, be neither or both.) That said, we do see the results strengthen slightly when we include an interaction for term length (table 2 column 7). And if we do try to just toss everything in together with interaction terms for chamber, term, and size, while significance craters, the regressors including term of office tend to remain statistically significant and closer to their original magnitudes while those based on size of chamber do not (table 2 column 8). This hint that a longer term of office leads to a greater focus on the public goods is food for a richer model.

The miniscule values of R-squared, never toping 2% 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 which results in deviation from the unitary-actor optimum. At this point it is worth investigating other sources of friction inhibiting the optimal map, especially the traditional redistricting principles highlighted by Winburn (2007).

To do so, we first calculate measures of 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 3, first row of coefficients). The average overlap between parent and offspring is 71.6% in both chambers. This, along with the low R-squared in table 2, is strong evidence that whatever constrains that shifts in partisan alignment is not a general constraint on moving district boundaries. The party in charge retains between 3 and 4 percentage points more of the previous district, consistent with the idea that incumbents do not want to lose voters with whom they have established a connection. While one would expect that more senior members would be even more reluctant to lose voters who have known them for such a long time, we do not find any interaction between this effect and seniority (not shown). We do see that the lower chamber places greater emphasis than the upper chamber on maximizing the overlap for seats currently held by incumbents (table 3, columns 1 and 2), consistent with our previous results that lower chambers display a greater bias toward the individual preferences of the legislators present in the negotiation. Interestingly, in the lower chamber where we found that competitive districts are helped more, there is no relationship between overlap and margin of victory, suggesting that these districts are helped without the need for greater displacement.

We then looked into whether traditional legal principles affect the extent to which districts are rearranged. Using data from Justin Levitt's redistricting website, we have coded which chambers are required to consider each of four standard requirements: compactness, respect for natural geography, respect for communities of interest, and minimizing splits of counties. Columns (5) and (6) of table 3 show the results for Upper and Lower chambers. In each case, laws requiring respect for compactness are an important limiter of the extent to which districts lines are changed. Not surprisingly, lower (upper) chambers with such a requirement exhibit 6.0 (7.7) percentage points greater overlap, presumably because the requirement limits the ability to reach out and "grab" the distant pockets of voters one wishes to graft onto an existing district. Moreover, the imposition of this restraint counters the overall tendency to keep one's own districts more intact than those of the opposition.

On the other hand, laws to consider communities of interest have essentially no effect, possibly because such definitions are vague and therefore difficult to litigate while the concept of compactness can be expressed by any number of mathematical formulae. As Winburn notes (p29), courts have weighed in relatively seldom on the principle of protecting communities of interest because the standards are ambiguous and because any area is characterized by communities along various dimensions—economic markets, racial and ethnic neighborhoods—that are often cross-cutting, implying that grouping by one dimension inevitably severs along another. While there is no single measure for compactness, they are highly correlated thus offering a justiciable principle.

Perhaps most interestingly, laws requiring respect for county lines and natural geography have clearly chamber-specific effects with the former having statistically significant effect only in upper chambers and the latter only in lower chambers. Surprisingly, requiring respect for county lines actually reduces the average overlap between parent and offspring districts and exacerbates the tendency to fiddle more with opposition districts. We speculate this is a result of the relative sizes of counties and districts. These laws typically require that the map attempt to minimize the number of counties split across districts. In a chamber with relatively few counties per district, counties will inevitably be split to maintain population equity. But in upper chambers, where there are more counties per district, these laws are constraints that prevent minor tinkering at the edges, requiring that whole counties be exchanged if district boundaries are to be shifted. Hence the result that in upper chambers, county restrictions actually reduce overlap and raise the difference between in-party districts, which the mapmaker tries to preserve the incumbent-voter connection and out-party districts, in which the mapmaker tries to sever that connection or at least pays it no mind. The role of natural geography, impinging more severely on the lower chambers, is perhaps the inverse case. With larger districts, one can always find a set of geographic boundaries to suit one's purpose while at the level of a finer partition, there are fewer relevant barriers and they thus become more dispositive.

All told, what we have is evidence that this form of moving friendly voters from safe districts to competitive districts so as to raise the expected probability of victory is:

(i) Clearly focused on districts already held by the party in charge of redistricting, consistent with our theory of broadly shared bargaining power and a status-quo threat.
(ii) Rare and modest, explaining a very small fraction of the overall change in competitive balance of a district
(iii) Possibly more acute when the term of office is shorter.

We also have evidence that districts held by the party in power are adjusted less overall, consistent with a benefit to the incumbent to keeping one's own voters.

V. 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 it 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 function of the competitiveness of the district. Why is there seemingly no systematic attempt to pick off the most vulnerable opposition districts as defined by the closeness of the most recent election? One possibility for this 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 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 but this would deepen the puzzle. Thus at present we do not have an answer.

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 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. We find support for this discontinuity and thus support for a process that privileges current members of the chamber.

References

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* vol 35, no 6, pp 878-904.

Cox, Gary W. and Jonathan Katz. 2002. *Elbridge Gerry's Salamander*. Cambridge University Press.

Daley, David. 2016 *Ratf**ked: Why Your Vote Doesn't Count*. WW Norton & Company, Inc.

Friedman, John N. and Richard T. Holden. 2008. Optimal Gerrymandering: Sometimes Pack, but Never Crack. *American Economic Review*, 98(1): 113-144.

Gul, Faruk and Wolfgang Pesendorfer. 2010. Strategic Redistricting. *American Economic Review* vol 100, pp 1616-1641.

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 and Stephen Ansolabehere, Stanford Election Atlas.

Winburn, Jonathan. 2007. *The Realities of Redistricting*. Lexington Books.

Table 1: Regression Discontinuity

	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
RD_estimate	0.01922	0.01681	0.02534	0.01954	0.02769	0.02933	0.01382	0.01382
Conventional SE	0.00951	0.00859	0.01234	0.01042	0.01659	0.01516	0.00684	0.00684
Convention p-value	0.043	0.051	0.04	0.061	0.095	0.053	0.043	0.043
Robust p-value	0.059	0.069	0.041	0.077	0.124	0.065	0.062	0.049
Kernel Type	Triangular							
BW Type	mserd						manual	
Observations	1588	1828	1588	1828	1588	1828	1588	1588
Order Loc. Poly. (p)	1	1	2	2	4	4	1	1
Order Bias (q)	2	2	3	3	5	5	2	2
BW Loc. Poly. (h)	0.061	0.072	0.07	0.108	0.094	0.118	0.12	0.12
BW Bias (b)	0.098	0.12	0.101	0.157	0.118	0.15	0.12	0.17
limit on the absolute value of the electoral margin	0.2	0.25	0.2	0.25	0.2	0.25	0.2	0.2

Table 2: Cross-section Regressions

Dependent variable: shift in normal democratic vote toward party drawing the lines

Units: 1 percentage point = 0.01

Method: OLS

Is the chamber competitive (seat share 33% - 67%)?	all	no	yes	all
Chamber	both	both	both	upper
	[1]	[2]	[3]	[4]
Constant	-0.00553	-0.01249**	0.00715	0.00763
	(0.00514)	(0.00585)	(0.00978)	(0.01560)
Own seat (held by party drawing the lines)	0.01546**	0.0192**	0.00632	-0.00325
	(0.00665)	(0.00803)	(0.01170)	(0.01890)
Margin of victory for party drawing lines	0.04763	0.0601	0.0356	-0.109
	(0.04310)	(0.04780)	(0.08660)	(0.12200)
size (log of number of seats)				
Long term (indicator of 4 year term)				
Own seat * Margin of Victory	-0.0947*	-0.08713	-0.105	0.087
	(0.05290)	(0.06190)	(0.09900)	(0.14700)
Own seat * size				
Own seat * longterm				
margin of victory * size				
margin of victory * long term				
Own party * margin of victory * size				
Own party * margin of victory * long term				
N	1875	1046	829	534
R squared	0.003	0.007	0.004	0.00200

Table 2 continued

Dependent variable: shift in normal democratic vote toward party drawing the lines				
Units: 1 percentage point = 0.01				
Method: OLS				
Is the chamber competitive (seat share 33% - 67%)?	all	all	all	all
Chamber	lower	both	both	both
	[5]	[6]	[7]	[8]
Constant	-0.00973**	0.0611	-0.00944	0.0445
	(0.00447)	(0.03900)	(0.00609)	(0.04720)
Own seat (held by party drawing the lines)	0.02234***	-0.1022**	0.03128***	-0.02984
	(0.00593)	(0.04760)	(0.00807)	(0.05770)
Margin of victory for party drawing lines	0.106***	-0.535	0.1043**	-0.1259
	(0.03850)	(0.34030)	(0.05160)	(0.43700)
size (log of number of seats)		-0.01474*		-0.01137
		(0.00854)		(0.00991)
Long term (indicator of 4 year term)			0.01236	0.004749
			(0.01120)	(0.01300)
Own seat * Margin of Victory	-0.1658***			0.1854
	(0.04750)			(0.49900)
Own seat * size		0.02623**		0.0129
		(0.01040)		(0.01210)
Own seat * longterm			-0.0420***	-0.03319**
			(0.01420)	(0.01640)
margin of victory * size		0.1297*		0.04868
		(0.07530)		(0.09240)
margin of victory * long term			-0.1811**	-0.1525
			(0.09280)	(0.11400)
Own party * margin of victory * size		-0.1889**		-0.0803
		(0.08690)		(0.10500)
Own party * margin of victory * long term			0.2761**	0.2209*
			(0.11400)	(0.13700)
N	1341	1875	1875	1875
R squared	0.011	0.007	0.016	0.017

Table 3: Overlap

Dependent variable: overlap between parent district and matched offspring district

Method of Estimation: OLS

sample restriction on overlap	none	none	> 50%	> 50%	none	none
Chamber	Upper	Lower	Upper	Lower	Upper	Lower
	[1]	[2]	[3]	[4]	[5]	[6]
Own seat (held by party drawing the lines)	0.0305**	0.0434***	0.0185*	0.0301***	0.0467*	0.0620***
	(0.0134)	(0.0089)	(0.0106)	(0.0070)	(0.0277)	(0.0177)
Margin of victory for party drawing lines	0.107***	0.00177	0.0596**	-0.0340**	0.0772**	-0.0347
	(0.0329)	(0.0224)	(0.0258)	(0.0173)	(0.0332)	(0.0228)
Compactness					0.0773***	0.0598***
					(0.0231)	(0.0153)
Own seat * Compactness					-0.0831***	-0.0158
					(0.0287)	(0.0192)
Retain Communities of Interest					0.00486	0.00956
					(0.0243)	(0.0170)
Own seat * Communities					0.0559*	0.0360*
					(0.0303)	(0.0217)
Natural Geography					0.0178	0.0658)***
					(0.0293)	(0.0209)
Own seat * Natural Geography					-0.0242	-0.0835***
					(0.0366)	(0.0261)
County lines					-0.0599**	0.00791
					(0.0263)	(0.0174)
Own seat * County lines					0.0653**	-0.0209
					(0.0328)	(0.0222)
Constant	0.716***	0.716***	0.771***	0.792***	0.685***	0.668***
	(0.0138)	(0.0092)	(0.0111)	(0.00724)	(0.0235)	(0.0148)
N	855	2319	786	2004	855	2319
R squared	0.018	0.010	0.010	0.010	0.058	0.036

Figure 1: Preserving takes precedence over flipping

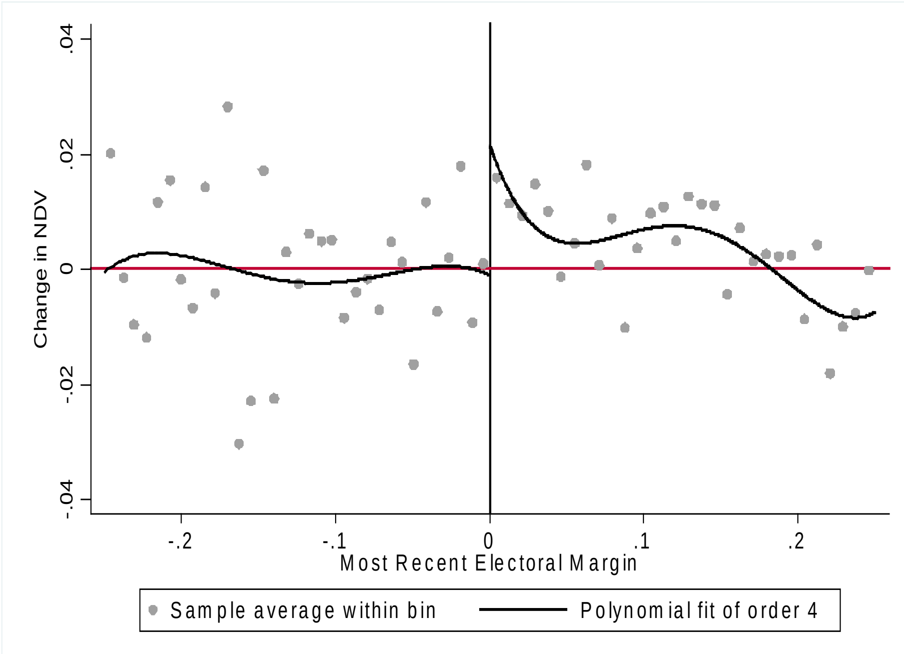

Figure 2: Compactness requirements inhibit evisceration