American Economic Review 100 (September 2010): 1616–1641
http://www.aeaweb.org/articles.php?doi=10.1257/aer.100.4.1616

Strategic Redistricting

By Faruk Gul and Wolfgang Pesendorfer*

Two parties choose redistricting plans to maximize their probability of winning a majority in the House of Representatives. In the unique equilibrium, parties maximally segregate their opponents' supporters but pool their own supporters into uniform districts. Ceteris paribus, the stronger party segregates more than the weaker one, and the election outcome is biased in the stronger party's favor and against the party whose supporters are easier to identify. We incorporate policy choice into our redistricting game and find that when one party controls redistricting, the equilibrium policy is biased towards the preferences of the redistricting party's supporters. (*JEL* C72, D72)

Regional differences in population growth periodically necessitate changing congressional districts' boundaries. This *redistricting* process creates an intense conflict between political parties. In this paper, we analyze the redistricting process for the House of Representatives and the interaction between redistricting and policy choice.

States face few constraints when setting their congressional districts' boundaries: congressional districts must have the same population as each other and must be contiguous, which, in practice, is a fairly permissive constraint. In a well-known example, Illinois's fourth congressional district combines two disjoint areas through a very narrow strip. In short, a political party that controls a state's political institutions has wide latitude in designing a favorable electoral map. In some cases, independent commissions rather than individual parties control the redistricting process. We ignore such bipartisan redistricting and assume that a single party controls each state's political institutions. Bipartisan redistricting can be incorporated into our model by giving parties control of less than 100 percent of the districts and interpreting the remainder as an exogenous nonpartisan redistricting plan.

We model the strategic interaction between the two parties as a zero-sum game under uncertainty. We recognize that parties and different agents within parties may evaluate election outcomes in different ways; incumbents may want to protect their own seats, while other party members may wish to maximize the number of representatives. We focus, however, on the most important election outcome: majority control in the House of Representatives.[1]

Our model combines Downs-Hotelling style party competition with redistricting. We assume that the two parties' supporters have different distributions of policy preferences (i.e., ideal points). The overall distribution of ideal points is a θ-weighted average of these two distributions where θ is the proportion of Republicans in the population. In Section II, we fix the party policies and study redistricting in isolation. Voter characteristics and an uncertain aggregate state determine the fraction of Republican voters in the population. Parties observe voter characteristics

* Gul: Princeton University, Department of Economics, Princeton, NJ 08544 (e-mail: fgul@princeton.edu); Pesendorfer: Princeton University, Department of Economics, Princeton, NJ 08544 (e-mail: pesendor@princeton.edu). We thank Thomas Eisenbach, David Epstein, Jonathan Katz, John Kim, Jay Lu, Thomas Palfrey, and Andrea Prat for helpful comments on earlier drafts of this paper. This research was supported by grants from the National Science Foundation (SES-0550540, SES-0518753).

[1] In Section III, we analyze how a desire to protect incumbents changes our model.

but do not know the aggregate state when they redistrict. We show that equilibrium is unique. In equilibrium, parties maximize the number of seats they would get if aggregate uncertainty were to resolve in a manner that yields each party half of the seats. We call this particular outcome of aggregate uncertainty the *critical state*. In the optimal redistricting plan, party 1 (the Republican Party) picks a cutoff characteristic, maximally segregates voters below the cutoff and combines all voters above the cutoff into uniform districts. Hence, parties segregate voters with unfavorable characteristics and combine voters with favorable characteristics. This description of the optimal strategy generalizes Guillermo Owen and Bernard Grofman's (1988) well-known *pack-and-crack gerrymandering strategy*.[2]

We compare the equilibrium behaviors of the *strong* and *weak* parties: assume, for simplicity, that the two parties face the same, symmetric ex ante distribution of characteristics. One party—the strong party—controls a larger *territory* than the other—the weak party. We show that the strong party will choose a more segregated redistricting plan. Specifically, the strong party will create fewer and more lopsided favorable districts. The weak party will create fewer unfavorable districts and more balanced favorable districts.

One focus of the empirical literature on redistricting (see, for example, Andrew Gelman and Gary King 1990; Gary W. Cox and Jonathan N. Katz 1999) is the notion of *bias*. These papers estimate a vote-seat curve that relates a party's vote share to its share of seats and define bias as the excess seat share that the party would have had with a vote share of ½ (i.e., seat share minus ½). Hence, the bias favors a party if its excess seat share is positive.[3] We show that the election is biased in the strong party's favor. Furthermore, in each territory, there is a *local bias* favoring the redistricting party. The local bias is always greater in the strong party's territory than in the weak party's territory. Thus, overall bias is related to local bias: if the election is biased in party 1's favor, then 1's territory will exhibit more bias than 2's territory.[4]

Cox and Katz (2002) study the evolution of local bias after Republican and Democratic redistricting plans between 1946 and 1970. This period encompasses the redistricting revolution (triggered by *Baker versus Carr* (1962) and subsequent Supreme Court decisions) which the authors argue greatly strengthened the Democratic Party.[5] Their results indicate that the prerevolutionary Republican redistricting plans' biases were larger than the postrevolutionary Republican redistricting plans' biases, while the opposite holds for Democratic redistricting plans. This finding is consistent with our model's predictions since the Republican Party was the stronger party before the revolution and became the weaker party afterwards.

Jonathan Rodden (2007) provides evidence that in the United States and other industrialized countries, left-leaning voters tend to be more concentrated than right-leaning voters. For example, in the 2000 presidential election, the smallest Democratic vote share in any congressional district was 24 percent while there were 24 districts with a Democratic vote share over 80 percent and five districts with a Democratic vote share over 90 percent. Our model predicts that the redistricting party will try to "pack" as many of its opponents as possible into designated losing districts.

[2] If the only constraint on redistricting is that the mean characteristic must equal some fixed constant, the strategy that maximizes the expected number of seats yields two types of districts, as in the work of Owen and Grofman. Note however that, in our model, a party's objective is to maximize the probability of winning a majority. When there is aggregate uncertainty, maximizing this objective function implies maximizing the expected number of seats at the critical state that would yield both parties the same number of seats.

[3] We use a slightly different measure of bias. Given an estimated vote-seat curve f, our measure would yield $0.5 - f^{-1}(0.5)$ while the empirical literature defines it as $f(0.5) - 0.5$. Since f is increasing, these two terms have the same sign; that is, qualitatively, the two notions are identical. Our definition is more convenient for our analysis because the state in which the two parties win an equal number of Representatives (rather than an equal number of votes) plays a key role in our model.

[4] This result assumes homogenous populations and symmetric ex ante voter characteristic distributions.

[5] *Baker versus Carr*, 369. U.S. 186(1962).

Hence, the 2000 election suggests that segregating Democratic voters was easier than segregating the Republican voters. We examine how such asymmetries affect the parties' electoral prospects. For example, suppose 2's supporters are easier to identify and therefore easier to segregate than 1's supporters. This could be due to the geographic concentration of 2's supporters or because 2's support is correlated with some observable variables such as ethnicity. We show that, if parties are otherwise in a symmetric situation, then the election will be biased in 1's favor. Identifying its own supporters is less valuable for a party than identifying its opponent's supporters since the optimal redistricting plan requires segregating opponent's supporters and pooling the party's own supporters into uniform districts.

Finally, we incorporate policy choice into our redistricting game and examine how redistricting plans affect this choice. Our extended model combines Downs-Hotelling style party competition with redistricting. For simplicity, we assume that one party controls all redistricting and show that the policy choice will be biased away from the overall median towards the policy preferences of the redistricting party's supporters. Equilibrium policy choice targets the median in the redistricting party's favorable districts because, when the election is close, those districts will be the most hotly contested. Notice that this result holds even though parties care only about winning a majority in the House of Representatives.

I. Related Literature

Our work builds on Owen and Grofman (1988), who show that when there is local uncertainty but no aggregate uncertainty, both the redistricting plan that maximizes the expected number of seats and the plan that maximizes the probability of winning a majority create two types of districts, ones that overwhelmingly favor the opponent and others that the party is expected to win.[6]

In our model, the redistricting party observes a one-dimensional signal that determines the probability a particular voter is a Democrat or a Republican. A key feature of our setting is that parties have limited information about voters' policy preferences. John N. Friedman and Richard T. Holden (2008a) show that when parties receive a sufficiently informative signal of each voter's policy preference, the optimal redistricting plan matches extreme types, i.e., the most favorable types are placed in the same district as the least favorable types, the second most favorable types are placed with the second least favorable types and so on. Friedman and Holden (2008b) analyze competitive redistricting. Here, the authors first extend their earlier paper's characterization of optimal strategies to the strategic setting. Later, they explore some consequences of abandoning the assumption that parties have near-perfect information about each voter's policy preferences. In Section IIIE below, we discuss their work in more detail and relate their information assumptions to the assumptions of Theorem 1.

Thomas W. Gilligan and John G. Matsusaka (2006) and Stephen Coate and Brian Knight (2007) study socially optimal redistricting plans. An important redistricting constraint is the mandate to create and maintain districts with a substantial majority of minority voters.[7] Charles Cameron, David Epstein, and Sharyn O'Halloran (1996), Epstein and O'Halloran (1999), Kenneth Shotts (2001), and Delia Grigg and Katz (2005) analyze the impact of majority-minority districts on the welfare of minorities. In our model, the mandate to create majority-minority districts amounts to a *lower bound* on segregation. Section IIIC below describes how this constraint can be incorporated into our model and briefly discusses its impact.

[6] For a different generalization of Owen and Grofman see Katerina Sherstyuk (1998).

[7] The Voting Rights Act mandates adequate representation of minorities, which courts have interpreted as a mandate to create districts with a significant majority of minority voters.

Cox and Katz (2002) provide a comprehensive study of redistricting since the reapportionment revolution of the 1960s. Their model (and much of the literature on redistricting) focuses on the trade-off between bias and responsiveness. There is also a large empirical literature that focuses on the so-called seat-vote curve that is generated by various redistricting plans (see, for example, King and Robert X. Browning 1987, Gelman and King 1990 and 1994).

Shotts (2002) and Timothy Besley and Ian Preston (2007) model the interaction of redistricting and policy choice. In Shotts' model, parties are policy motivated and redistrict to move the median Representative closer to the party ideal point. Besley and Preston examine how partisan bias affects a party's willingness to accommodate swing voters. In their model, parties have policy preferences, but swing voters constrain their extremism. The partisan bias of the electoral map affects this constraint and, hence, affects policy. The mechanism connecting policy and redistricting in our model is different: polarizing policies facilitate segregation, and therefore strong parties polarize despite not having any policy preference.

II. Voting

In this section, we introduce a stochastic median voter model with fixed party positions. This model provides the framework for the redistricting game we describe in Section III. Our model has the following key features: (i) parties redistrict based on incomplete information about voters' party affiliations, (ii) variation in the abilities of district candidates causes correlation in voter behavior within a district, and finally (iii) aggregate factors that affect the fortunes of the two parties cause correlation in voter behavior at the national level. The details of our model are described below.

Voters have symmetric, single-peaked preferences over the policy space $\mathbb{R}$. A voter's ideal point is unknown to parties at the time of redistricting. However, the parties can observe the Republican (and hence the Democratic) vote share in each voting precinct[8] and other demographic information. Based on this information, parties assess the probability that a particular voter is a Republican or a Democrat. Republicans have ideal points drawn from the cumulative distribution I_1, and Democrats have ideal points drawn from I_2.

After redistricting but before the election, some voters move or change their party affiliations. An *aggregate state* affects the transition of voter affiliations. To be concrete, let the mechanism through which the aggregate state affects voting behavior be as follows: each voter is replaced with probability 2α, and the replacement will be a Republican with probability s_0 if the aggregate state is s_0. Hence, if the share of Republican voters at the time of redistricting in a given large population is ω_0, the corresponding share on voting day will be $(1 - 2\alpha)\omega_0 + 2\alpha s_0$.

We call $\omega = (1 - 2\alpha)\omega_0 + \alpha$ the *characteristic* of this pool of voters and let $s = 2\alpha s_0 - \alpha$ be the (re-scaled) aggregate state. When the pool consists of all voters in a particular district, we call ω the district characteristic. Hence, if state s occurs, the ideal point distribution on the day of voting in a district with characteristic ω is

$$(\omega + s)I_1 + (1 - \omega - s)I_2. \tag{1}$$

Note that $\omega \in \Omega := [\alpha, 1 - \alpha]$ and $s \in S := [-\alpha, \alpha]$.

Each voter votes either for party 1 (Republican Party) or 2 (Democratic Party). The Republican policy is fixed at 1 and the Democratic policy is fixed at -1. Voters' preferences depend on their ideal points and on a noise term d. If $d < 0$, the voter is inclined towards the Republican candidate and if $d > 0$ the voter is inclined towards the Democratic candidate. We interpret d as

[8] A voting precinct is the smallest geographical area for which vote shares of parties are observable.

a valence term quantifying the competence difference between the two candidates in a district. Note that d is district specific while the aggregate state s affects all districts. A voter with ideal point x gets utility $u_1(x, d) = -|1 - x| - d$ if the Republican candidate is elected and $u_2(x, d) = -|-1 - x| + d$ if the Democratic Party candidate is elected. Therefore, this voter prefers party 1 if $u_1(x, d) > u_2(x, d)$; that is, if

$$(2) \qquad 2d < |-1 - x| - |1 - x|,$$

and party 2 if this inequality is reversed. For $x \in [-1, 1]$, the inequality above is equivalent to

$$(3) \qquad d < x.$$

The variable d can compensate for a less favorable policy; negative values indicate a voter's willingness to choose party 1 despite the fact that party 2 offers a policy closer to his ideal point. Let L be the cumulative distribution function of d. We assume:

(i) L is strictly concave on $\mathbb{R}_+$, continuous, and symmetric around 0, i.e., $L(d) = 1 - L(-d)$ for all d.

Republican voters have ideal points greater than zero while Democratic voters have ideal points less than zero. That is:

(ii) $I_1(0) = 0$; $I_2(0) = 1$.

Finally, we assume that:

(iii) I_1 is strictly increasing and convex on $[0,1]$, has median in $[0,1]$, and is continuous, and $I_2(x) = 1 - I_1(-x)$ for all $x \in [-1,0]$.

If I_1 has a density, then the curvature restriction says that this density is increasing on the interval $[0,1]$. For example, if I_1 has a single peaked density with median and mode equal to 1, then our assumption is satisfied. The symmetry assumption requires that the distribution of Democratic ideal points on the interval $[-1,0]$ be the mirror image of the distribution of Republican ideal points on the interval $[0,1]$.

The symmetry assumptions are made for convenience. Our main results, Theorems 1–3, are unaffected if we drop symmetry and impose a curvature assumption on I_2 analogous to the curvature assumption on I_1. By contrast, the curvature restrictions on I_1, I_2 and L are important for our analysis; they are needed for proving that the district outcome function is continuous and has the right curvature properties.[9]

Let $\theta = \omega + s$ be the proportion of Republicans in the district on voting day. The median for a given θ, $x(\theta)$, is the ideal point x that solves

$$(4) \qquad \theta I_1(x) + (1 - \theta) I_2(x) = \frac{1}{2}.$$

Increasing the proportion of Republicans moves $x(\theta)$ to the right. Assumption (iii) implies that, for each θ, there is a unique median and that this median is strictly increasing in θ with

[9] See Lemma 1 below.

$x(1/2) = 0$. Party 1 wins the district if $d < x(\theta)$ and, therefore, the probability that party 1 wins a district with a θ-proportion of Republicans is

$$\pi(\theta) := L(x(\theta)). \tag{5}$$

We call the function π the *District Outcome Function* (*DOF*). Lemma 1 below establishes that the DOF is concave over the range in which 1 is more likely to win and convex over the range in which 2 is more likely to win. Hence, as the leading party's support increases, its probability of winning increases at a decreasing rate.

LEMMA 1. *The district outcome function π is continuous, strictly concave on $[½, 1]$, and $\pi(\theta) = 1 - \pi(1 - \theta)$ for all $\theta \in [0, 1]$.*

PROOF:

See the online Appendix.

Assumptions (i)–(iii) provide a median voter model with local noise that yields a DOF with the curvature and symmetry described in the lemma above. Assumption (i) requires the distribution of local uncertainty to be S-shaped and continuous. If L admits a density, the assumption implies that this density has mode zero and is single peaked. We can interpret this assumption as saying that small differences in candidate competence are more likely than large differences.[10] Suppose the ability to avoid scandals is how candidate competence is measured. Then, our curvature assumption would be violated if, for example, there is either no scandal or a large scandal but no intermediate scandals.

The DOF has a particularly simple form when there is no local uncertainty. Suppose $L(d) = 1$ for $d > 0$ and $L(d) = 0$ for $d < 0$. Then, the DOF is π^∞ such that

$$\pi^\infty(\theta) = \begin{cases} 1 & \text{if } \theta > ½ \\ ½ & \text{if } \theta = ½ \\ 0 & \text{if } \theta < ½. \end{cases} \tag{6}$$

Notice that π^∞ is inconsistent with assumption (i) since it is not continuous. However, if L is normal with mean zero and variance $1/n$, then, for large n, the resulting π approximates π^∞ and satisfies assumptions (i)–(iii). Hence, we can study the no local uncertainty case as a limit.

III. Redistricting

In the preceding section, we have described a stochastic median voter model (with fixed policy positions). In the redistricting game, each party maximizes its probability of winning a majority in the House by optimally allocating voters among the districts it controls. To avoid integer problems and simplify the analysis, we assume that each party controls a continuum of districts. The mass of districts under 1's control is $\lambda \in [0, 1]$ and 2 controls the remaining mass $1 - \lambda$. The districts under the control of a party are called its *territory*. Unless otherwise stated, we will assume that both parties' territories have positive mass. When parties redistrict, they know voter characteristics. Parties allocate these characteristics (i.e., voters) among a continuum of

[10] We thank a referee for pointing out this implication of Assumption (i).

equal-sized districts; that is, they choose a district-characteristic distribution over their territory. For simplicity, we assume that the average characteristic in both territories is ½.[11]

Given these assumptions, a *redistricting plan H* is a cumulative distribution function (cdf) with mean ½ and support contained in Ω; $H(z)$ is the share of districts that have district characteristic no greater than z. Let $\mathcal{F}$ denote the collection of all such cdfs. The *segregation constraint* is the most dispersed (or segregated) feasible distribution of district characteristics. The cdf $F \in \mathcal{F}$ is the segregation constraint for party 1 and $G \in \mathcal{F}$ is the segregation constraint for party 2. A redistricting plan H is *feasible for party* 1 (*party* 2) if and only if F (G) is a mean preserving spread of H.[12] We write $H' \succeq H$ if H is a mean preserving spread of H'.

The segregation constraint arises from the fact that parties have limited information about voters. Suppose that there are two kinds of districts in party 1's territory—favorable and unfavorable—and that each favorable district had characteristic $r > ½$ while each unfavorable district had characteristic $1 - r < ½$. Suppose, further, that the party has no other information about voters. Thus, the voter's characteristic is either r (if he belongs to a favorable district) or $1 - r$ (if he belongs to an unfavorable district). In this case, it is not feasible for party 1 to create a district with characteristic lower than $1 - r$ or higher than r since it can identify only voters with characteristics r or $1 - r$. On the other hand, the party can combine groups of voters with these two characteristics to create any district characteristic between these values. One possible redistricting plan is to make all districts have characteristic ½. Another possibility is a uniform H on the interval $[r, 1 - r]$. Note that in either case, the original two-point distribution is a mean preserving spread of the redistricting plan. By combining voters from the two types of districts in appropriate proportions, the party can create any distribution H of district characteristics such that the original two-point distribution is a mean preserving spread of H. Conversely, creating a distribution that does not have this property would require more information than is available, and therefore such distributions are not feasible.

If the segregation constraint, F, has a single element in its support, then F itself is the only feasible redistricting plan. We rule out this trivial case and assume that both parties face a nondegenerate segregation constraint, i.e., $0 < F(z) < 1$ for some z. Since there is a continuum of districts, the "law of large numbers" ensures that, with strategy H, party 1 wins $D(H, s)$ districts in its territory in state s, where

$$D(H, s) = \int \pi(\omega + s)\, dH(\omega). \tag{7}$$

Party 1's total seat share in the House given state s and the strategy profile (H, H') is

$$\Delta(H, H', s) = \lambda D(H, s) + (1 - \lambda) D(H', s) \tag{8}$$

and therefore party 1 wins the election if $\Delta(H, H', s) \geq ½$. Hence, 1 chooses H to maximize

$$\Pr\{s \mid \Delta(H, H', s) \geq ½\}. \tag{9}$$

Party 2 chooses H' to minimize this probability. Parties do not know s and have beliefs on $S = [-\alpha, \alpha]$ for $\alpha \in [¼, ½)$. The cumulative distribution of these beliefs is strictly increasing on S and continuous. These assumptions ensure that (i) neither party can win the election with probability 1, (ii) there must exist an aggregate state, the *critical state*, at which both parties win the

[11] This assumption is made for convenience. The model easily generalizes to the case where the average voter characteristic is different in the two territories.

[12] The cdf F is a mean preserving spread of H if F, H have the same mean and $\int_{\alpha}^{x} [F(\omega) - H(\omega)]\, d\omega \geq 0$ for all x.

same number of representatives, and (iii) the probability of such a tie is zero. Theorem 1 below asserts that the equilibrium strategy profile is unique. This result requires the existence of such a critical state. Beyond guaranteeing the existence of critical state, the aggregate state distribution plays no role in our analysis, and therefore we do not specify it.

A *redistricting game* is a triple $\Lambda = (F, G, \lambda)$, where F, G are the redistricting constraints of 1 and 2 respectively and λ is the size of party 1's territory. If $G = F$, we sometimes write $\Lambda = (F, \lambda)$. We assume that parties choose their redistricting plans simultaneously. In practice, redistricting is done rather infrequently, and parties rarely choose their plans simultaneously. Our analysis is robust to the timing of moves: any sequencing of redistricting decisions would lead to the same equilibrium outcome as our simultaneous move game.

A. *An Example with No Local Uncertainty*

In this simple example, we assume no local uncertainty as described in equation (6) above. Hence,

$$\pi^{\infty}(\theta) = \begin{cases} 1 & \text{if } \theta > 1/2 \\ 1/2 & \text{if } \theta = 1/2 \\ 0 & \text{if } \theta < 1/2. \end{cases} \tag{10}$$

The segregation constraints F and G assign probability ½ to ¼ and probability ½ to ¾. Party 1 controls two-thirds of the districts, i.e., $\lambda = 2/3$.

Suppose that 1 establishes two kinds of districts in its territory: unfavorable districts "packed" solely with 2 supporters (i.e., $\omega = 1/4$) and mixed districts that 1 expects to win. Since 1 controls ⅔ of the electoral map, if the share of the favorable districts in its own territory is above ¾, then it wins more than $2/3 \times 3/4 = 1/2$ of all districts and therefore wins the election. Because party 1 has a larger territory than party 2, we can guess that with equal vote shares party 1 will win the election. This means that in the critical state party 2 will have more votes than party 1, which means that 1 will not win any district in 2's territory. Hence, in the critical state, party 1 must tie the election without winning any districts in 2's territory. That is, 1 must win at least ¾ of its own districts.

To create a ¾-proportion of favorable districts, 1 must combine all of 1's supporters with half of the 2 supporters. The average characteristic in these mixed districts will be

$$\omega = \frac{1}{3} \times \frac{1}{4} + \frac{2}{3} \times \frac{3}{4} = \frac{7}{12}. \tag{11}$$

Hence, 1 will win the election as long as $7/12 + s > 1/2$ and, therefore, party 1's equilibrium payoff is at least $\Pr\{s > -1/12\}$.

Note that there is no strategy for party 1 that enables it to win ¾ of the districts in its territory when $s < -1/12$. On the other hand, by creating uniform districts, 2 can ensure that it wins its entire territory whenever $s < -1/12$. Therefore, 2 can guarantee winning the election whenever $s < -1/12$ and hence its equilibrium payoff is no less than $\Pr\{s \leq -1/12\}$. It follows that 1's equilibrium payoff is $\Pr\{s > -1/12\}$ and that 2's equilibrium payoff is $\Pr\{s \leq -1/12\}$. Moreover, party 1's strategy described above and the uniform redistricting plan for 2 constitute an equilibrium. In equilibrium, both parties choose redistricting plans that maximize their seat shares at $s = -1/12$, i.e., the state in which the election is tied. It is easy to verify that 1's equilibrium strategy is unique. However, 2 can choose other redistricting plans and still win all districts in its territory

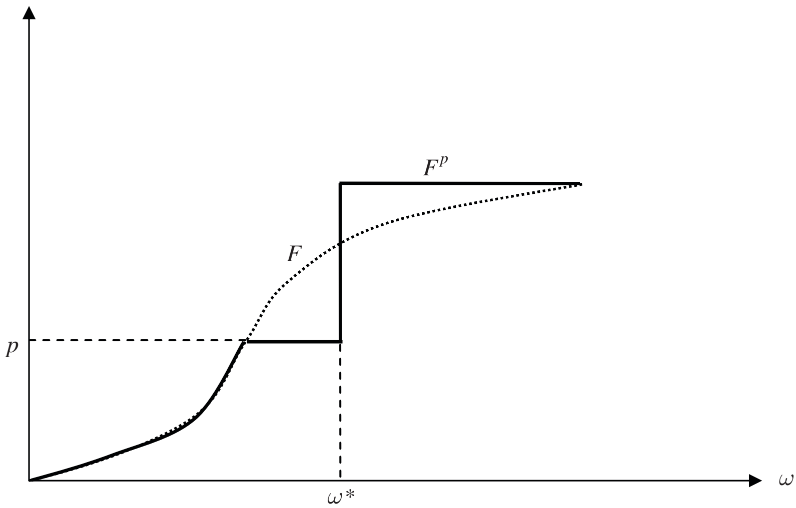

Figure 1. p-segregation

when in state $s = -\frac{1}{12}$. Hence, 2 has multiple equilibrium strategies in this example. In Theorem 1, π is strictly increasing (i.e., there is local uncertainty), and hence this multiplicity is ruled out.[13]

B. *Equilibrium Strategies*

We show in Theorem 1 below that party 1's equilibrium strategy fully segregates the lower p-percentile of characteristics and creates a mass of uniform districts with the same average from the upper $1 - p$-percentile. If F is continuous, then there is a cutoff z such that for $\omega \leq z$, the redistricting plan coincides with the segregation constraint, while characteristics above z are combined into uniform districts ω^*. We call such a strategy a *p-segregation plan* and let F^p denote the p-segregation plan for the redistricting constraint F. The Appendix provides the formal definition of p-segregation plans. In Figure 1, we illustrate a p-segregation plan for a continuous segregation constraint F.

Characteristics that favor party 1 are unfavorable for 2, and hence a p-segregation plan for party 2 fully segregates all districts above some critical ω and creates a mass of uniform districts with the same average below ω. With some abuse of notation, we write G^p for party 2's p-segregation plan given the segregation constraint G. Theorem 1 establishes that the equilibrium is unique and that equilibrium strategies are p-segregation strategies.

THEOREM 1. *There exist (i) p, q such that (F^p, G^q) is the unique equilibrium of Λ and (ii) a unique s^* that solves $\Delta(F^p, G^q, s) = \frac{1}{2}$. In equilibrium, parties maximize their seat shares in state s^*.*

[13] The example can be viewed as the limit case as local uncertainty disappears. Say that local uncertainty disappears along the sequence Λ^n if π^n, the DOF of Λ^n, converges pointwise to π^∞. Then, it is easy to show that in the limit equilibrium (i) at the critical state, the stronger party wins half of all districts despite not winning any districts in the opponent's territory, and (ii) the weaker party chooses a uniform redistricting plan.

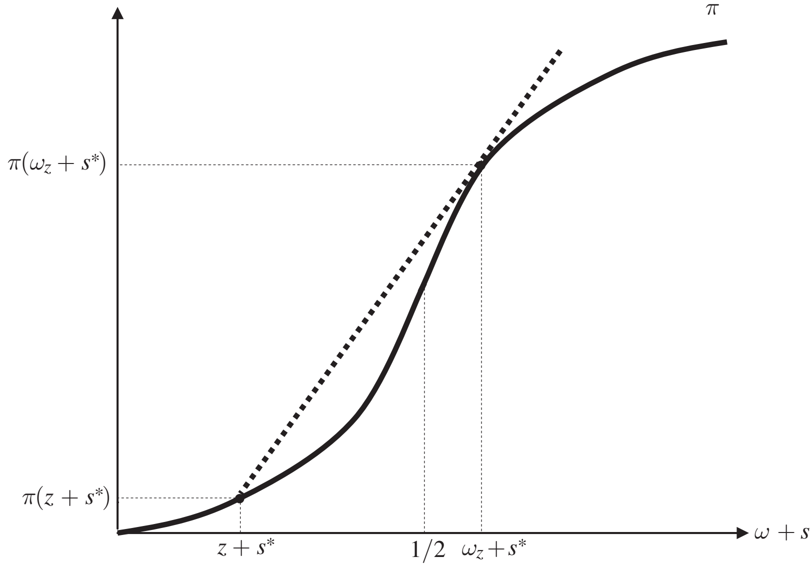

Figure 2. Optimal Strategy

Theorem 1 shows that a single parameter characterizes a party's optimal strategy. Henceforth, we identify equilibrium strategies with the pair (p, q). We call the state at which the election is tied in the unique equilibrium the *critical state* and denote it $s(\Lambda)$. Parties' redistricting plans maximize their seat shares in the critical state. To see why, assume party 1 can improve its seat share in the critical state by choosing some nonequilibrium strategy. Then, continuity implies that the party also wins in states slightly below (worse than) the critical state, and hence party 1 wins a majority with greater probability, contradicting the optimality of the equilibrium strategy.

The p-segregation plan maximally segregates districts in the convex part of π (districts that are most favorable to the opponent). We can provide a simple characterization of party 1's optimal p if F has a density $f > 0$ and π is differentiable. Let $s^* = s(\Lambda)$ be the critical state, $p = F(z)$. Define $\omega_z = \big(1/(1-p)\big)\int_{\omega \geq z} \omega f(\omega)\, d\omega$. Hence, ω_z is the average characteristic in the favorable districts given the distribution F. Then, 1's optimal strategy must satisfy:

$$\pi'(\omega_z + s^*)(\omega_z - z) = \pi(\omega_z + s^*) - \pi(z + s^*). \tag{12}$$

We illustrate this equation in Figure 2.

To understand this equation, suppose party 1 considers marginally raising the cutoff z (and hence raising p). Raising z increases the average characteristic in the favorable districts. This benefit yields an increase $(1-p)\big(\partial\pi(\omega_z + s^*)/\partial z\big) = f(z)\pi'(\omega_z + s^*)(\omega_z - z)$ in the mass of districts that party 1 wins. Raising z also decreases the proportion of districts with the favorable characteristic. The cost of this decrease is $-f(z)\big(\pi(\omega_z + s^*) - \pi(z + s^*)\big)$. At the optimum these two effects must cancel.

C. *Incumbency Protection and Majority-Minority Districts*

We assume that parties maximize the probability of winning a majority in the House of Representatives. As we argue in the introduction, given the rules of the House of Representatives, this is the natural party objective. However, individual representatives also have an incentive to guarantee their own election, and therefore parties may have to balance incumbency protection with the objective of achieving a majority.

To incorporate the incumbency protection objective into our model, suppose an ι-fraction of districts in party 1's territory are incumbent districts that must be protected. In particular, party 1 decides that incumbents should carry their districts with a probability of at least $\gamma > ½$. If this goal is feasible, then there is a district characteristic $\hat{\omega}$ such that

$$\gamma = E\pi(\hat{\omega} + s) \tag{13}$$

$$\iota \leq 1 - F(\hat{\omega}) \tag{14}$$

(where the expectation is over the aggregate uncertainty s).

For simplicity, assume that the redistricting constraint F has support $\{\alpha, 1 - \alpha\}$. Let F^p be the optimal redistricting plan for party 1 ignoring the need for incumbency protection, and let ω^* denote the district characteristic in the favorable districts in F^p's support. If $\omega^* \geq \hat{\omega}$ and $\iota \leq 1 - p$ (i.e., the optimal favorable district meets the requirement for incumbency protection and there are enough districts to protect all incumbents), then incumbency protection does not constrain the party. The party simply ensures that each incumbent has a favorable ω^*-district.

If one of these constraints is violated, i.e., there are too few favorable districts in the optimal redistricting plan $(\iota > 1 - p)$ or the favorable districts are not safe enough $(\omega^* < \hat{\omega})$, then the (constrained) optimal redistricting plan can be found in two steps: first, create an ι-proportion of $\hat{\omega}$-districts. Then, compute the new redistricting constraint, $\hat{F}$, for the remaining districts.[14] Then, choose an optimal redistricting plan, as described in the proof of Theorem 1, given the new constraint $\hat{F}$. Hence, parties will still choose p-redistricting plans, but $\iota \times \lambda$-fraction of districts will be beyond the control of either party.

Clearly, there is a trade-off between the need to protect incumbents (as measured by γ) and the party's electoral success. A more sophisticated objective function would be one where parties trade off the probability of winning a majority against incumbency protection. While an analysis of this trade-off is beyond the scope of this paper, it is clear that the equilibrium redistricting strategy that emerges from such a model is as described above: a fraction of districts will be designed to protect incumbents, and a p-segregation plan will be applied to the remainder. [15]

Like incumbency protection, the Voting Rights Act, which courts have interpreted as a mandate to create districts with a significant majority of minority voters, constrains parties' redistricting options. If minority voters favor the Democratic Party, then the Voting Rights Act means that there will be an ι fraction of districts with average characteristic at or below $\hat{\omega} < ½$. Assume that meeting the mandate is feasible for either party. Then, for the Democratic Party, the mandate's effect is similar to that of protecting incumbents: if $\hat{\omega} < \omega^*$ or if $\iota > 1 - p$, the mandate constrains the Democratic Party, and the optimal redistricting plan is determined by the two-step procedure outlined above.

For the Republican Party, creating districts that are sufficiently favorable to the opponent can never be a constraint since the optimal redistricting plan does just that. However, the optimal redistricting plan may not create enough opponent-favorable districts to meet the mandate. Consider again the case in which F has a binary support $\{\alpha, 1 - \alpha\}$. Since the mandate is feasible,

[14] District $\hat{\omega}$ has a δ proportion of voters with characteristic α and a $1 - \delta$ proportion of $1 - \alpha$'s where $\delta\alpha + (1 - \delta) \times (1 - \alpha) = \hat{\omega}$. Subtracting from the original distribution $\delta\iota$ mass of α's and $(1 - \delta)\iota$ mass of $1 - \alpha$'s and normalizing (i.e., dividing by $1 - \iota$) yields $\hat{F}$.

[15] Michael C. Herron and Alan E. Wiseman (2008) use observed redistricting plans to assess redistrictors' intentions. Their focus is on theories of legislative policy choice. An extended version of our model that includes incumbency protection could be used to infer a party's weight on incumbency protection from the observed redistricting plans.

it must be that $\alpha \leq \hat{\omega}$. If $p \geq \iota$, then the mandate does not constrain party 1; it simply ensures that an ι fraction of the districts favorable to party 2 are minority districts. If $p < \iota$ however, then the optimal plan for party 1 does not yield enough party 2-favorable districts to meet the mandate. In that case, the constrained optimal redistricting plan is $\hat{p} = \iota$.[16]

D. *Multiple Redistricting Constraints*

We have assumed that a single cumulative distribution function describes all redistricting constraints that a party faces within its territory. This would be the case, for example, if a party could draw its electoral map without regard to state boundaries and other geographical considerations. In practice, parties are not allowed to create districts that cross state lines or connect distant regions within a state. Theorem 1 can easily be extended to deal with these additional restrictions. We can model a party's territory as a collection of *regions* and define a different redistricting constraint for each region. In equilibrium, each party will choose a distinct p-redistricting plan for every region in its territory. Each such plan will maximize the party's seat share at the critical s.

We assume a continuum of districts. This essential simplifying assumption enables us to apply the law of large numbers to the district level (local) uncertainty. Since there are 435 congressional districts, this aspect of the assumption seems relatively unproblematic. The continuum assumption also allows us to ignore integer constraints. In reality, the feasible distributions of district characteristics can only approximate the optimal p-segregation plan since there are finitely many districts in each territory. The p-segregation plans are accurate approximations of the optimal plan for large territories (such as California or New York) but less accurate approximations for small territories.

E. *Comparison to the Friedman and Holden Analysis*

Friedman and Holden (2008a) develop and study a model with a single party facing an optimal redistricting problem. Friedman and Holden (2008b) extend the analysis of the earlier paper to a competitive two-party setting as in our model. Below, we illustrate the relationship between their model and our model with a simple example that has some features of both.

There are N districts, and party 1 (referred to simply as "the party" below) controls all of them. The party observes a signal, $y \in [-1, 1]$, about each voter's ideal point, $x \in [-1, 1]$. The cdf $H(\cdot\,|\,y)$ describes the distribution of a voter's ideal point given that his signal is y. Assuming that there are many voters in each district, an informal appeal to the law of large numbers ensures that the posterior distribution of ideal points in a district with signal distribution F is $G_F(x) = \int H(x\,|\,y)\; dF(y)$.

The probability of winning a district is a function of the post-redistricting median in that district. A higher median implies a higher probability of winning. Let $m(G)$ be the median of G and let F^0 be the distribution of signals. The goal of the party is to create N districts, i.e., cumulative distributions of y's, $F^1, \ldots, F^N$ so as to maximize

$$\sum_{i=1}^{N} W\big(m(G_{F^i})\big) \qquad \text{subject to} \qquad \frac{1}{N}\sum_{i=1}^{N} F^i \;=\; F^0. \tag{15}$$

[16] We are grateful to a referee for pointing out this effect. The analysis above is simplified by the assumption of a binary support for F. In the general case, there are many ways to create a district with average characteristic $\hat{\omega}$. As a result, the details of the segregation constraint affect the constrained-optimal redistricting plan, making its computation more involved. However, the conclusion stays the same–we can describe the optimal redistricting plan by a two-step procedure in which first the constraint is met and then the optimal p-redistricting plan for an appropriately defined residual redistricting constraint is found.

The difference between the optimal redistricting plans in Friedman and Holden and in our paper stems largely from different (simplifying) assumptions about the conditional distributions $H(\cdot\,|\,y)$. Friedman-Holden assume that signals provide near-perfect information of a voter's ideal point.[17] That is, $H(\cdot\,|\,y)$ is concentrated around y for each y. Under the Friedman-Holden assumption, the optimal redistricting plan "matches extremes," that is, pools the worst signals with the best signals and the next worst signals with the next best signals, and so on.

There are two important features of the Friedman-Holden optimal redistricting plans: first, and most important, for every N, the distribution of medians associated with their optimal plan stochastically dominates any other feasible distribution. Hence, Friedman-Holden optimal plans are optimal *for any increasing* W. Therefore, the Friedman-Holden characterization of optimal plans is valid even in a strategic setting. Second, optimal plans never segregate extreme unfavorable types or create identical districts of favorable types.

In contrast, we assume that $H(\cdot\,|\,y) = \big((1 + y)/2\big)\, H(\cdot\,|\,1) + \big((1 - y)/2\big)\, H(\cdot\,|\,-1)$. For example,

$$H(x\,|\,y) = \frac{1 + x - y(1 - |x|)}{2} \tag{16}$$

satisfies our requirements. In that case, signal -1 implies ideal points uniform on $[-1, 0]$, signal 0 implies ideal points uniform on $[-1, 1]$, and signal 1 implies ideal points uniform on $[0, 1]$. In addition, we make the following curvature assumption: consider the effect of replacing a small mass of y-type voters with the same mass of y'-type voters $(y' > y)$ on $W\big(m(G_F)\big)$. We assume this effect is S-shaped in $m(G_F)$, i.e., the benefit increases as $m(G_F)$ increases up to a critical value and then decreases. Under our assumptions, an optimal redistricting plan for the N districts is a finite analogue of the p-segregation strategy described in Theorem 1 and converges to a p-segregation strategy as N goes to infinity.

Notice that our assumptions imply that signals cannot be very informative.[18] Therefore, the Friedman-Holden paper and our paper are complementary. Together, they illustrate how optimal redistricting depends on the available information. Our results apply when information is coarse, while the Friedman and Holden characterization applies when parties have precise information about voter ideal points.

IV. Party Strength, Bias, and Segregation

In this section, we relate the redistricting game's parameters to equilibrium outcomes. Theorem 2 below shows that, for comparative statics analysis, keeping track of the critical state, $s(\Lambda)$, is sufficient. Specifically, we show that, if the parameters of the redistricting game change so that party i wins over a greater range of aggregate states (the party becomes stronger), then i chooses a larger p. For $\hat{p} > p$, $F^{\hat{p}}$ is a mean preserving spread of F^p since the two distribution functions have the same mean and satisfy the standard single crossing property for distributions (Peter A. Diamond and Joseph E. Stiglitz 1974). Hence, when party i wins over a greater range of states, it chooses a more segregated redistricting plan with more lopsided favorable districts and a larger proportion of maximally segregated unfavorable districts.

[17] Near-perfect signals is a sufficient condition. Friedman and Holden (2008a) give a nice example that does not have near-perfect information but leads to their characterization of optimal plans. Similarly, our conditions are sufficient conditions, and there are examples that violate them but lead to redistricting plans as characterized in Theorem 1 above.

[18] The only case where the two models overlap is when there are two signals $y \in \{0, 1\}$. In that case, the signal can be near perfect and satisfy our assumption as well.

We say that a parameter change in the redistricting game Λ makes party 1 *stronger* if the critical state, $s(\Lambda)$, falls. Similarly, a parameter change makes party 2 stronger if $s(\Lambda)$ rises. The following theorem establishes that the stronger a party gets, the more it segregates.

THEOREM 2. *Let* $\Lambda = (F, G, \lambda)$, $\hat{\Lambda} = (F, \hat{G}, \hat{\lambda})$ *and let* $p, \hat{p}$ *be the corresponding equilibrium strategies of* 1. *If* 1 *is stronger in* Λ *than in* $\hat{\Lambda}$, *then* $p \geq \hat{p}$.

PROOF:

See the Appendix.

Using Figure 2 from Section III, we can provide a straightforward intuition for Theorem 2. Let $\hat{z}$ be the optimal cutoff when the critical state is $s(\hat{\Lambda})$, and let z be the optimal cutoff when the critical state is $s(\Lambda)$. Since $s(\hat{\Lambda}) > s(\Lambda)$ and ω_z is increasing in z, the tangency illustrated in the figure implies that $\hat{z} < z$.[19]

In Theorem 2, party 1's redistricting constraint is fixed, while the other parameters of the game may change. Although the distribution of aggregate uncertainty plays no role in our analysis, it does affect a party's probability of winning. If the probability distribution over states remains constant, then increasing a party's strength increases its probability of winning. However, Theorem 2 remains valid even if the probability distribution over states changes as other parameters change. In that case, a party's strength refers to its ability to win in unfavorable circumstances and not to its probability of winning.

We say that the game, Λ, is biased in 1's favor if $s(\Lambda) < 0$ and in 2's favor if $s(\Lambda) > 0$. If $s(\Lambda) < 0$ (> 0), then party 1 (2) can win the election even though a majority of voters prefer party 2 (1). The bias in territory i is defined analogously. Let

$$s_1(\Lambda) := \{s \mid D(F^p, s) = \tfrac{1}{2}\} \tag{17}$$

$$s_2(\Lambda) := \{s \mid D(G^q, s) = \tfrac{1}{2}\} \tag{18}$$

where (p, q) is the unique equilibrium of Λ. Hence, $s_i(\Lambda)$ is the aggregate state that would yield a tie in territory i. Arguments analogous to the ones made for $s(\Lambda)$ ensure that $s_i(\Lambda)$ is also well defined.

Theorem 3 below establishes that the local bias always favors the redistricting party. Also, it shows that the local bias increases when the redistricting party becomes stronger. Finally, Theorem 3 shows that bias grows as the strong party gets stronger.

THEOREM 3. (*i*) *For any* Λ, $s_1(\Lambda) \leq 0 \leq s_2(\Lambda)$ *and* $s_1(\Lambda) \leq s(\Lambda) \leq s_2(\Lambda)$. (*ii*) *Let* $\Lambda = (F, G, \lambda)$, $\hat{\Lambda} = (F, \hat{G}, \hat{\lambda})$. *If* $s(\Lambda) \leq s(\hat{\Lambda})$, *then* $s_1(\Lambda) \leq s_1(\hat{\Lambda})$. (*iii*) *The critical state* $s(F, G, \lambda)$ *is decreasing in* λ.

PROOF:

See the Appendix.

Theorem 3 relies on two key observations: let p_s be the p that maximizes party 1's seat share, $D(F^p, s)$, in state s. In Theorem 2, we showed that the stronger party 1 is the more it segregates;

[19] We are grateful to a referee for suggesting use of the figure to explain Theorem 2.

that is, p_s is decreasing in s. The second observation is that, fixing s, as p increases towards its optimal level, the seat share increases; that is, $D(F^p, s)$ is increasing at $p < p_s$.

Let $s = s(\Lambda)$ and $s_1 = s_1(\Lambda)$. First, assume that $s < 0$. In that case, party 2 can win more than half of the seats in its territory in state s (for example, a uniform redistricting plan would yield more than half of the seats to 2). Since the election is tied at s, party 1 must win more than half of the seats in its territory as well. Hence, we have

$$D(F^p, s) \geq \Delta(F^p, G^q, s) = \frac{1}{2} = D(F^p, s_1). \tag{19}$$

Then, the monotonicity of D implies that $s_1 \leq s < 0$.

Next, assume $s \geq 0$ and therefore $p_s \leq p_0$ and

$$D(F^{p_0}, 0) \geq D(F^{p_s}, 0) \geq D(F^0, 0) = \frac{1}{2}. \tag{20}$$

The last equality follows since, at $s = 0$, a uniform redistricting plan yields each party exactly half the seats. It follows that $D(F^{p_s}, 0) \geq$ ½ and $s_1 \leq s$. Parts (ii) and (iii) follow from similar arguments.

Theorems 2 and 3 offer testable implications of our model. Increasing party 1's strength increases the local bias in territory 1. Cox and Katz (1999) provide evidence on the evolution of bias after the Republican and Democratic redistricting from 1946 to 1970. This period encompasses the redistricting revolution (triggered by Supreme Court decisions starting with Baker versus Carr 1962) which the authors argue greatly strengthened the Democratic Party in the sense defined above. Their results (Table 3, p. 830) indicate that the prerevolution Republican redistricting plans yielded larger biases than postrevolutionary Republican redistricting plans, while the evolution of the biases is exactly reversed for Democratic redistricting plans. Cox and Katz define bias as the difference between the seat share of a party and ½ when its vote share is one half. They estimate that the bias of Republican plans drops from 8.26 percent to 0.092 percent, while the bias of Democratic plans increases from 4.76 percent to 8.70 percent. We can use their estimates to compute the estimated bias according to the definition used here.[20] In that case, the estimated bias for Republican plans drops from 2.3 percent to essentially zero, while the estimated bias for Democratic plans increases from 1.1 percent to 2.1 percent.

The following corollary summarizes our comparative statics results under the assumption that parties face the same redistricting constraint but have different sized territories. For any distribution $F \in \mathcal{F}$, let $\rho(F)$ denote the distribution of $1 - \omega$. The distribution F is *symmetric* if $\rho(F) = F$. We say that the game is *symmetric* if both parties face the same symmetric redistricting constraint F. Hence, in a symmetric redistricting game both parties' situations are identical except for the sizes of their territories. The following corollary shows that, if the game is symmetric, the election will be biased in favor of the party with the larger territory; the stronger party will choose a more segregating redistricting plan (i.e., create more lopsided districts) and generate a more biased electoral map in its territory.

COROLLARY 1. *If Λ is symmetric and $\lambda >$ ½, then the election is biased in* 1*'s favor;* 1 *segregates more than* 2 *and enjoys a greater local bias than* 2.

The comparative statics results above provide some insight into how equilibrium redistricting plans differ from ex post seat maximizing redistricting plans. Suppose a particular state $s > s(\Lambda)$

[20] Computing the biases defined above from their estimated seat-vote curve is straightforward.

is realized, and 1 wins the election. Since 1's redistricting plan maximizes its seat share at $s(\Lambda)$ but not at s—that is, the optimal redistricting plan at s has less segregation (smaller p) than the equilibrium plan—1 will win many districts with larger margins of victory than would be optimal in the seat maximizing plan. Hence, it may appear as if 1 is creating overly safe districts. In contrast, 2's redistricting plan will appear as if it has segregated too little; its seat share would increase had it created more safe districts. A symmetric argument applies for $s < s(\Lambda)$. Thus, the winner will appear to be overly conservative while the loser will seem overly aggressive in its redistricting.

We conclude this section with an analysis of how differences in segregation opportunities affect the equilibrium outcome. A mean preserving spread of the segregation constraint means that the party is less constrained, and therefore, it is stronger. Such a change may come about through better information, that is, greater ability to identify voters. Note that $F = G$ does not mean that both parties have the same segregation opportunities. For example, if one party has some supporters that are easily identified by their ethnicity or their address while the other party has no such reliable indicators of support but both parties face the same distribution of voters, we will have $F = G$, but F will be different at low values than at high values.

Note, however, that if $F = G$ and F is symmetric (i.e., $\rho(F) = F$), then both parties' supporters are equally difficult to segregate. More generally, suppose F is not symmetric. Let $H \in \mathcal{F}$ be such that H coincides with F for $\omega < 1/2$ and is symmetric. Hence, H is the symmetric characteristic distribution in which both parties' supporters are distributed like party 2's supporters in F. If H is a mean preserving spread of F, then we can conclude that, in the game with $F = G$, party 2's supporters are more "spread out" than party 1's supporters and therefore easier to segregate.

DEFINITION 1: *Party 2's supporters are easier to segregate at F if there is $H \in \mathcal{F}$ such that $\rho(H) = H$, $H(\omega) = F(\omega)$ for $\omega < 1/2$ and $F \succeq H$.*

Example: Let $\Omega = \{3/8, 1/2, 9/16\}$ and assume F puts probability 0.25 on $3/8$, probability 0.25 on $1/2$, and probability 0.5 on $9/16$. In this case, party 2's supporters are easier to segregate because the symmetric distribution H that puts probability 0.25 on $3/8$, 0.5 on $1/2$, and 0.25 on $5/8$ is a mean preserving spread of F.

Examining US election results, in particular, the outcomes in the two parties' safe districts suggests that they face different segregation constraints. In the 2000 presidential election, the smallest Democratic vote share in any congressional district was 24 percent, while there were 24 districts with Democratic vote shares over 80 percent and 5 districts with a Democratic vote share over 90 percent. This suggests that there are stronger indicators of Democratic voting proclivities than of Republican voting proclivities.[21]

Theorem 4 examines a situation in which both parties have equal-sized territories, with the same characteristic distribution. If $F = \rho(F)$, then both parties face the same constraint and hence

$$s(\Lambda) = 0. \tag{21}$$

Hence, the redistricting game is *unbiased*, i.e., the party with majority support wins the election. When 2's supporters are easier to segregate, the critical state is less than 0 and the election is biased in party 1's favor.

[21] See Rodden (2007) for further evidence that left-leaning districts tend to be more lopsided than right-leaning districts.

THEOREM 4. *If party 2's supporters are easier to segregate in $\Lambda = (F, \frac{1}{2})$, then $s(\Lambda) \leq 0$.*

Theorem 4 establishes that the equilibrium outcome is *biased against* the party whose supporters can be segregated more readily. To understand this result, consider a change that increases both parties' ability to segregate party 2's supporters: this change does not help party 2 in territory 2 because its equilibrium strategy (the *p*-segregation plan) creates uniform districts of supporters. However, since maximally segregating the opponent's supporters is optimal, party 1 benefits from its increased ability to segregate party 2's supporters.

Theorem 4 can be strengthened to establish a strict inequality $(s(\Lambda) < 0)$ if the extreme supporters of 2 are more extreme than the extreme supporters of 1. More formally, let $\underline{\omega}(F)$ be the minimum element in the support of F (the strongest supporter of 2) and let $\overline{\omega}(F)$ be the maximum element in the support of F (the strongest supporter of 1). If $\overline{\omega}(F) < 1 - \underline{\omega}(F)$, then 1 strictly gains from its greater ability to segregate 2's supporters.

V. Policy Choice

So far, we have assumed that parties have different and fixed policy positions and therefore attract different kinds of voters. That the two parties competing for Congress in fact hold distinct and fairly stable policy positions seems uncontroversial. Identifying the source of this differentiation is beyond the scope of this paper. Instead, we ask a more limited question: suppose parties can vary their policy positions only on some dimensions; how does redistricting alter parties' policy choice?

We consider the voting model from Section II with a new policy dimension. Let y_i be party i's policy in this new dimension. Party positions on the original policy dimension are 1 and -1 as in Section II. In this extended model, a voter with ideal point x has utility

$$u_1(x, y_1, d) = -|1 - x| - \beta|y_1 - x| - d \tag{22}$$

from electing the party 1 Representative and utility

$$u_2(x, y_2, d) = -|-1 - x| - \beta|y_2 - x| + d \tag{23}$$

from electing the party 2 Representative. The parameter $\beta \in [0, 1]$ is a measure of the relative importance of the new policy dimension.[22] The utility function in Section II is a special case of the one above with $\beta = 0$.

We can interpret our model as a setting in which parties have limited abilities to commit. The fixed policy dimension represents policy choices for which commitment is not possible. In that case, voters substitute the ideal point of the party to predict the future policy choice. The new dimension represents a policy choice where parties' announcements prior to the election are credible, and hence parties are able to commit.

Party 1 wins the district if

$$u_1(x(\theta), y_1, d) > u_2(x(\theta), y_2, d). \tag{24}$$

[22] Since $\beta \leq 1$, we are assuming that the new policy dimension is no more important than the old dimension.

Let $\pi_y(\theta)$ be the probability that 1 wins a district with θ-share of Republicans given policies $y = (y_1, y_2)$. Let $d(\theta)$ be the value of d that solves

$$u_1(x(\theta), y_1, d) = u_2(x(\theta), y_2, d). \tag{25}$$

That is,

$$d(\theta) = x(\theta) + \frac{\beta}{2}(|y_2 - x(\theta)| - |y_1 - x(\theta)|), \tag{26}$$

Then, define

$$\pi_y(\theta) := L(d(\theta)). \tag{27}$$

Note that whenever $y_1 = y_2$, we have $d(\theta) = x(\theta)$, and therefore π_y is the same DOF as the one studied in Section II. Hence,

$$\pi_y(\theta) = \pi(\theta) = L(x(\theta)) = L(d(\theta)) \tag{28}$$

whenever $y_1 = y_2$. We maintain Assumptions (i)–(iii) from Section II, and therefore π_y has the same properties as the function π in the previous sections when $y_1 = y_2$.

The policy choices affect the competence differential required to win a district and, hence, affect the probability of winning that district. Expression (26) implies that choosing a policy as close as possible to the district median maximizes the probability of winning. Of course, different districts have different median ideal points, and therefore, in general, no policy is optimal in every district. Moreover, the aggregate state affects the district median, and therefore increasing the probability of winning in some aggregate state may reduce the probability of winning in other states.

To examine the interaction between redistricting and policy choice, we consider the simple case in which party 1 controls all districts ($\lambda = 1$). Hence, party 1's redistricting constraint, F, parameterizes the redistricting game. Party 1 first chooses its redistricting plan and its policy, and then party 2 chooses its policy. We employ the sequential structure to rule out mixed strategy equilibria. Pure strategy equilibria when party 1 first chooses a redistricting plan and then both parties simultaneously choose policies would be identical to the equilibria analyzed below. However, we cannot ensure the existence of a pure strategy equilibrium for the latter game.[23]

As in the previous sections, a feasible redistricting plan H is such that the redistricting constraint F is a mean preserving spread of H. In aggregate state s, party 1 wins a seat share of

$$D_y(H, s) = \int \pi_y(\omega + s)\, dH(\omega) \tag{29}$$

and, hence, wins the election if $D_y(H, s) > ½$. Party 2 wins the election if this inequality is reversed.

[23] Since party 1 controls redistricting it is natural to think of party 1 as the incumbent party. In that case, it seems plausible that party 1 would have to choose its policy first, and party 2, the challenger, can move second.

Theorem 5 below shows that the equilibrium outcome of the redistricting-policy game is unique. In that outcome, party 1 chooses the p-redistricting plan that would have been optimal for $\beta = 0$, that is, in the situation without the new policy dimension. Parties choose identical policies that cater to the party 1–favorable districts.

More precisely, let F^p be the optimal redistricting plan, and let s denote the critical state. In that redistricting plan, party 1 creates a mass of identical favorable districts. Let $\omega^*(p)$ be the common characteristic of these favorable districts. In the critical state s, the ideal point of the median voters in those districts is $x(\omega^*(p) + s)$. In the equilibrium, both parties choose the policy $x(\omega^*(p) + s)$.

THEOREM 5. *Assume* $\beta \leq 1$. *The unique equilibrium outcome of the redistricting-policy game is the redistricting plan* F^p *and the policies* $y_1 = y_2 = x(\omega^*(p) + s) > 0$.

PROOF:
See the Appendix.

To gain intuition for Theorem 5, note that the policy choice, like the redistricting plan, must maximize the probability of winning a majority of Representatives. In particular, this implies that both choices must maximize the seat share in the critical state s. There are two types of districts, districts that favor party 1 and districts that favor party 2. In the critical state, any policy to the right of the median in party 1–favorable districts $\big(x(\omega^*(p) + s)\big)$ cannot maximize the seat share since it is to the right of the median in every district.

Note that a policy to the left of $x(\omega^*(p) + s^*)$ cannot be optimal: recall that a basic property of optimal redistricting plans is that, in the critical state, districts favorable to party 1 (the party in charge of redistricting) are less lopsided than districts favorable to party 2. Now, suppose both parties choose the policy $x(\omega^*(p) + s)$ and consider a leftward deviation by one of the parties $\big(y_i < x(\omega^*(p) + s)\big)$. This will increase the deviator's probability of picking up a seat in districts favorable to party 2 but reduce the probability of picking up a seat in districts favorable to party 1. Note that the impact of a policy change is greater in districts that are more closely contested. Hence, in districts favorable to party 1, the negative impact of the leftward shift is greater than the corresponding positive impact in districts favorable to party 2. Moreover, since the election is tied in the critical state, it follows that there are more districts favorable to party 1 than districts favorable to party 2. As a result, the leftward shift in policy reduces the deviator's seat share in the critical state and therefore reduces the deviator's payoff.[24]

To evaluate the interaction of redistricting and policy choice, consider the benchmark case in which all districts are identical and have type $\omega = ½$. Then, the median preferred policy in the critical state is the optimal policy. It is straightforward to verify that the critical state is 0 in this case, and therefore the median voter's ideal point in every district is zero.

The equilibrium policy when party 1 redistricts differs from this benchmark for two reasons. First, parties cater to districts that favor party 1, that is, districts with type $\omega^*(p) > ½$. Second, the fact that party 1 has an advantage through redistricting implies that party 1 can win in aggregate states that are favorable to party 2. This effect shifts the critical state to the left, i.e., $s < 0$. The first effect $(\omega^*(p) > ½)$ moves the equilibrium policy to the right, while the second effect $(s < 0)$ mitigates the rightward shift. Overall, the policy tilts towards the right since $\omega^*(p) + s > ½$ and therefore $x(\omega^*(p) + s) > 0$. Hence, the effect of catering to the districts favorable to party 1 outweighs the effect of the leftward shift in the critical state.

[24] The intuitive argument implicitly assumes that the policy deviation is not too large so that districts favorable to party 2 remain the more lopsided districts. The assumption $\beta \leq 1$ is required to deal with large deviations.

When local uncertainty is small, party 1 wins a district with near certainty even if it has a very small advantage. Therefore, $\omega^*(p) + s$ is close to ½ in the optimal redistricting plan. When $\omega^*(p) + s$ is close to ½ then the median voter's ideal point is close to zero in the critical state, i.e., $x(\omega^*(p) + s)$ is close to zero. As a result, the equilibrium policies in the redistricting-policy game are close to the benchmark policies.

In the limit when there is no local uncertainty, the two effects described above exactly offset each other, and gerrymandering is policy neutral. The party in control of redistricting wins more often but still chooses the same policy as in the benchmark case. To make this precise, let L^n be a normal distribution with mean zero and standard deviation $1/n$. As discussed in Section II, L^n converges to L^∞ (no local uncertainty) as n goes to infinity. Let $y = (y_n^1, y_n^2)$ be the equilibrium policy choices for the redistricting-policy game with local uncertainty L^n. We have the following corollary:

COROLLARY 2. *Let* (H_n, y_n^1, y_n^2) *be the unique equilibrium in the redistricting-policy game with local uncertainty* L^n. *Then,* $\lim_{n\to\infty} y_n^1 = 0$.

Theorem 5 shows that policies are biased towards the party that controls redistricting because policies will target districts favorable to that party. However, the advantage gained from redistricting means moves the critical state to the left, and this effect mitigates the impact of redistricting on policy outcomes. In particular, when local uncertainty is small, gerrymandering affects only the probability of winning and has almost no effect on the policy choices.

VI. Conclusion

We have described how aggregate uncertainty creates strategic interaction between parties' redistricting decisions. This uncertainty ensures that one party's redistricting plan affects the other party's optimal action even though the fraction of districts a party wins at any particular state s is a separable function of its own and its opponents redistricting plans. Despite the vital role of aggregate uncertainty, the distribution of this uncertainty does not affect equilibrium strategies. It follows that asymmetric information regarding this distribution will have no effect on equilibrium outcomes.

Our model provides a framework for analyzing the interaction between redistricting and other decisions. We have considered one such interaction by adding a policy choice to our model. Other decisions such as the allocation of campaign resources across districts or the policy choices of individual candidates who care only about the outcome in their own district can also be studied within our framework.

Mathematical Appendix

A. *p-Segregation Plans*

In this subsection, we provide formal definitions of F^p and G^p. For any cdf H, let $\mu(H)$ be the mean of H and, for $p \in [0,1)$, let H_+^p be the distribution of the upper $1 - p$-percentile, i.e.,

$$H_+^p(\omega) := \max\left\{\frac{H(\omega) - p}{1 - p}, 0\right\}.$$

Let $H_+^1 := H$.

DEFINITION 2: *For any distribution F, let* $F^1 = F$ *and* F^0 *be the distribution that yields* $\mu(F)$ *for sure. For* $p \in (0, 1)$, *define* F^p *as follows:*

$$F^p(\omega) = \begin{cases} F(\omega) & \text{if } F(\omega) < p \\ p & \text{if } p \leq F(\omega),\ \omega < \mu(F_+^p) \\ 1 & \text{if } \mu(F_+^p) \leq \omega \end{cases}$$

for $p \in (0,1)$. *We call* F^p *the p-segregation plan for party* 1.

Characteristics that are favorable for party 1 are unfavorable for 2. Hence, the analog of F^p for party 2 fully segregates all districts above some critical ω and creates a mass of uniform districts with the same average below ω. For any distribution H, let $\rho(H)$ denote the corresponding distribution of $1-\omega$. That is, $\rho(H)$ is the unique distribution such that $\rho(H)(\omega) = 1 - H(1-\omega)$ at every continuity point of $\rho(H)$.[25] If G is the segregation constraint for 2, then $\rho(G)$ is the translation of G that makes G comparable to F, the segregation constraint of 1. If $\rho(G) = F$, then both parties face the same segregation constraint.

DEFINITION 3: *The p-segregation plan for* 2 *is the distribution* $\rho[\rho(G)^p]$. *With some abuse of notation, we let* G^p *denote* 2*'s p-segregation plan.*

B. *Preliminary Results*

LEMMA 2. *For* $F, G \in \mathcal{F}$, *both* $D(F, G, \cdot)$ *and* $\Delta(F, G, \cdot)$ *are continuous and increasing functions.*

PROOF:

Since $\theta = \omega + s$, both θ and π are increasing and continuous in s. Then, equations (7) and (8) in Section III yield the continuity and increasingness of $D(F, G, \cdot)$ and $\Delta(F, G, \cdot)$ respectively.

Define, for $\theta \in [0, ½)$ and $\theta' \in [½, 1]$,

$$g(\theta, \theta') = \frac{\pi(\theta') - \pi(\theta)}{\theta' - \theta}.$$

Recall that π is strictly concave on $[½, 1]$ and therefore there is a unique θ' that maximizes $g(\theta, \cdot)$ for every $\theta \in [0, ½)$. For all such θ, let $\phi(\theta)$ be this maximizer. By the Theorem of the Maximum, ϕ is continuous. The symmetry of π around ½ and the strict concavity of π on $[½, 1]$ imply that $½ < \phi(\theta) < 1 - \theta$ for all $\theta \in [0, ½)$ and that ϕ is (strictly) decreasing. It follows that $\lim_{\theta \to 1/2} \phi(\theta) = ½$. Hence, we can extend ϕ continuously to the closed interval $[0, ½]$ and note that the resulting ϕ is also decreasing with $\phi(½) = ½$. We record these observations in the lemma below:

LEMMA ϕ. *The function* ϕ *is continuous and decreasing. For all* $\theta \in [0, ½)$, $½ < \phi(\theta) < 1 - \theta$ *and* $\phi(½) = ½$.

[25] Note that $\rho(\rho(H)) = H$ and hence $\rho^{-1} = \rho$.

For any $F \in \mathcal{F}$, let F^p_- be the distribution of the lower p-percentile. That is,

$$F^p_- = \rho\big(\rho(F)^{1-p}_+\big).$$

Recall that F^p_+ is the distribution of the upper $1-p$-percentile and therefore

$$F = (1-p)F^p_+ + pF^p_-.$$

Let $\underline{\omega}(F)$ $(\overline{\omega}(F))$ be the minimum (maximum) of the support of F, and let $F^-(\omega)$ be the left limit of F at ω. For $p \in [0, 1]$, let $\omega^*(p, F) = \mu(F^p_+)$ and

$$\omega(p, F) = \{\omega \in [\underline{\omega}(F), \overline{\omega}(F)] \,|\, F^-(\omega) \le p \le F(\omega)\}$$

$$W_s(p, F) = \{\omega^*(p, F) + s - \phi(\omega + s) \,|\, \omega \in \omega(p, F)\}.$$

When the choices of F and s are clear, we write $\omega(p)$, $\omega^*(p)$, and $W(p)$ instead.

A correspondence h from the reals to nonempty subsets of the reals is increasing if $x \ge x'$, $w \in h(x)$, $w' \in h(x')$ imply $w \ge w'$ and is strictly increasing if the second inequality above is strict whenever the first one is strict.

LEMMA 3. *W and $\omega(\cdot)$ are increasing in p and both are upper-hemicontinuous* (*uhc*).

PROOF:

That $\omega(\cdot)$ is uhc and increasing is obvious. Note that $\omega^*(p)$ is the expectation of F conditional on a realization in the top $1-p$-percentile. Hence, it is a continuous and increasing function. It then follows from Lemma ϕ that the correspondence $-\phi\big(\omega(\cdot, F)\big)$ is increasing, W is increasing, and both are uhc.

The next three lemmas characterize the seat maximizing redistricting plan. Below, we omit the reference to party 1 and a p-segregation plan is always a p-segregation plan for party 1. Fix an aggregate state s and define the following maximization problem:

$$\max D(\hat{F}, s) \text{ subject to } \hat{F} \succeq F. \tag{A3}$$

LEMMA 4. *For every $F \in \mathcal{F}$, the maximization problem* (A3) *has a unique solution F^* such that $F^* = F^p$ for some $p \in [0, 1]$.*

PROOF:

See online Appendix.

In the proof of Lemma 4, we suppressed the dependence of p_F and ω_F on s. Below, we analyze how these variables change as s changes while F stays fixed. To simplify the notation we now suppress the dependence of p and ω on F. By Lemma 4, there exists a function $p : S \to [0, 1]$ such that F^{p_s} is the unique solution to

$$\max D(\hat{F}, s) \text{ subject to } \hat{F} \succeq F.$$

Henceforth, we write p_s to denote the solution to the above maximization problem. We write ω_s to denote the corresponding ω_F (as defined in the proof of Lemma 4) given state s. The following lemma shows that p_s is decreasing in s.

LEMMA 5. *The value p_s is a decreasing function of s; it is strictly decreasing at s such that $p_s > 0$.*

PROOF:

See online Appendix.

LEMMA 6. *If $\hat{p} < p \leq p_s$, then $D(F^{\hat{p}}, s) \leq D(F^p, s)$.*

PROOF:

See online Appendix

Symmetric arguments establish that there exists q_s such that $G^{q_s} := \rho[(\rho(G))^{q_s}]$ minimizes $D(H, s)$ among all $H \succeq G$, that q_s is increasing in s, and that $D(G^{q_s}, s) \leq D(G^q, s) \leq D(G^{\hat{q}}, s)$ whenever $\hat{q} < q \leq q_s$. Henceforth, it will be understood that Lemmas 4, 5, and 6 also entail the analogous statements for party 2.

C. *Proofs of Theorems* 1, 2, *and* 3

PROOF OF THEOREM 1:

By the Theorem of the Maximum p_s, q_s and $D(F^{p_s}, s)$, $D(G^{q_s}, s)$ are continuous functions of s and hence so is $\Delta(F^{p_s}, G^{q_s}, s)$. Note that $\Delta(F^{p_s}, G^{q_s}, s) \leq 1/2$ at $s = \underline{s}$ and $\Delta(F^{p_s}, G^{q_s}, s) \geq 1/2$ at $s = \overline{s}$. Hence, there exists s^* such that $\Delta(F^{p_{s^*}}, G^{q_{s^*}}, s^*) = 1/2$. To complete the proof we will show that $(F^{p_{s^*}}, G^{q_{s^*}}, s^*)$ is the unique equilibrium.

Since $\Delta(H, \hat{H}, s)$ is strictly increasing in s for all H, $\hat{H}$ (Lemma 2), party 1's payoff is greater than $\Pr\{s > s^*\}$ if and only if $\Delta(H, \hat{H}, s^*) \geq 1/2$. The strategy $F^{p_{s^*}}$ is the unique strategy that ensures $\Delta(F^{p_{s^*}}, \hat{H}, s^*) \geq 1/2$ for all feasible $\hat{H}$. But since this is a zero-sum game, it follows that $F^{p_{s^*}}$ is the unique equilibrium strategy for party 1. Symmetric arguments establish that $G^{q_{s^*}}$ is the unique equilibrium strategy for party 2.

PROOF OF THEOREM 2:

The proof follows immediately from Theorem 1 and Lemma 5.

PROOF OF THEOREM 3:

The proof of part (i) was presented in the discussion following the statement of Theorem 3.

Part (ii): Let $s = s(\Lambda)$, $s_1 = s_1(\Lambda)$, $\hat{s} = s(\hat{\Lambda})$, and $\hat{s}_1 = s_1(\hat{\Lambda})$. Part (i) implies $\hat{s} \geq s \geq s_1$ and hence, Lemma 5 implies $p_{\hat{s}} \leq p_s \leq p_{s_1}$. Lemma 6 then implies

$$D(F^{p_{\hat{s}}}, s_1) \leq D(F^{p_s}, s_1) \leq D(F^{p_{s_1}}, s_1).$$

By definition, $D(F^{p_s}, s_1) = 1/2$ and therefore $D(F^{p_{\hat{s}}}, s_1) \leq 1/2 = D(F^{p_{\hat{s}}}, \hat{s}_1)$. Then, $\hat{s}_1 \geq s_1$ as required.

Part (iii): Since $s_1 \leq s \leq s_2$, party i's seat share at s in territory i is weakly greater than ½. This implies that the critical state $s(F,G, \lambda)$ is weakly decreasing in λ.

D. *Proof of Theorems* 4 *and* 5

LEMMA 7. *If F^p maximizes $D(H, s)$ subject to $F \succeq H$ for $F \in \mathcal{F}$ and $D(F^p, s) \geq 1/2$, then $p < 1/2$.*

PROOF:

If $p > 0$ then, by Lemma 4, $\omega^*(p) + s = \phi(\omega_s + s)$, and by Lemma ϕ, $½ < \phi(\omega_s + s) < 1 - \omega_s - s$. Note that

$$D(F^p, s) \leq p\pi(\omega_s + s) + (1-p)\pi(\omega^*(p) + s).$$

Hence, we have $p\pi(\omega_s + s) + (1-p)\pi(\omega^*(p) + s) < p\pi(\omega_s + s) + (1-p)\pi(1 - \omega_s - s)$ and therefore, $D(F^p, s) < p\pi(\omega_s + s) + (1-p)\pi(1 - \omega_s - s)$. Then, the symmetry of π yields $p < ½$ as desired.

PROOF OF THEOREM 4:

First, we demonstrate that at the critical state s, p_s, $q_s < ½$. If $s \geq 0$ then $D(F^{p_s}, s) \geq D(F^{p0}, s) \geq ½$. If $s \leq 0$ then $D(G^{q_s}, s) \leq ½$ and, since $\lambda D(F^{p_s}, s) + (1-\lambda)D(G^{q_s}, s) = ½$, it again follows that $D(F^{p_s}, s) \geq ½$. We can therefore apply Lemma 7 to conclude that $p_s < ½$. A symmetric argument establishes $q_s < ½$.

Since $q_s < ½$, $p_s < ½$, it follows that the plans $F^{½}$, $G^{½}$ are mean preserving spreads of the equilibrium strategies.

Consider the strategy $[\rho(F)]^{½}$. Since party 2's supporters are easier to segregate at F, it follows that $[\rho(F)]^{½}$ is a feasible strategy for party 1, and, since $[\rho(F)]^{½}$ is a mean preserving spread of $[\rho(F)]^q$, it follows that $[\rho(F)]^q$ is a feasible strategy for party 1. Clearly, if party 1 chooses $[\rho(F)]^q$ and party 2 chooses $\rho([\rho(F)]^q)$, then $s(\Lambda) = 0$ since the strategies are identical. We conclude that $s(F, \lambda) \leq 0$.

PROOF OF THEOREM 5:

Let F^p be the seat maximizing redistricting plan when both parties choose identical policies and s be the corresponding critical state. First, we show that $p > 0$ and $s < 0$. If $p = 0$, then $s = 0$ and $\omega^*(p) + s = ½$. But since F is non-degenerate, $\underline{\omega} + s < ½$ and $\phi(\underline{\omega} + s) > ½$ by Lemma ϕ. The argument given in the proof of Lemma 4 then implies that $p = 0$ is not an optimal strategy. The uniqueness of the optimal redistricting plan then implies that $s < 0$.

Since $p > 0$, it follows that $\omega^*(p) + s > ½$ and therefore $x(\omega^*(p) + s) > 0$ as required. Finally, let $\theta^* := \omega^*(p) + s$ and let $\theta := \omega + s$ where ω is any element in the support of F^p such that $\omega \neq \omega^*(p)$. Note that $\omega \leq \omega_s$. Let $y_1 = y_2 = x(\theta^*)$ and $y = (y_1, y_2)$.

Claim: *Let $y' = (y_1', y_2)$ and $y'' = (y_1, y_2'')$. Then, (i) $y_1' \neq y_1$ implies $D_{y'}(F^p, s^*) < ½$; (ii) $y_2'' \neq y_2$ implies $D_{y''}(F^p, s^*) > ½$.*

PROOF:

Part (i): By definition, $D_y(F^p, s^*) = ½$. Since $y_1' > y_1$ reduces party 1's probability of winning in every district, it follows that $D_{y'}(F^p, s^*) < D_y(F^p, s^*)$ if $y_1' > y_1$. Consider therefore $y_1' < y_1 = y_2 = x(\theta^*)$. The probability of winning district θ^* is $L(x(\theta^*) - \beta(x(\theta^*) - y_1'))$, and the probability of winning district θ is bounded above by

$$\min\{L(x(\theta) + \beta x(\theta^*) - \beta y_1'), L(x(\theta) + \beta x(\theta^*) - \beta x(\theta))\}$$

since $x(\theta) < 0$. Since $p < ½$, it follows that the proportion of districts with type θ^* is greater than ½. Therefore, it suffices to show that the loss in district type θ^* from choosing y_1' instead of $y_1 = x(\theta^*)$ is greater or equal to the maximal gain in any district with type θ. That is:

$$L(x(\theta^*)) - L\big(x(\theta^*) - \beta(x(\theta^*) - y_1')\big) \geq \min\big\{L\big(x(\theta) + \beta(x(\theta^*) - y_1'\big), L\big(x(\theta) + \beta(x(\theta^*) - x(\theta))\big)\big\} - L(x(\theta)). \tag{A6}$$

To prove inequality (A6), note that since $\theta \leq \omega_s + s$, Lemma ϕ implies that $1 - \theta > \theta^* > ½$ and therefore we have $x(\theta^*) > 0 > x(\theta)$ and $x(\theta^*) < -x(\theta)$. Since L is S-shaped and symmetric around 0 it follows that for all $\Delta \in [0, x(\theta^*) - x(\theta)]$

$$L(x(\theta^*)) - L(x(\theta^*) - \Delta) \geq L(x(\theta) + \Delta) - L(x(\theta)).$$

Since $\beta \leq 1$ this in turn implies that inequality (A6) holds. This proves part (i).

Part (ii) trivially holds if $y_2'' > y_2 = x(\theta^*)$ since this reduces the probability of winning in all districts. Hence, consider $y_2'' < x(\theta^*)$. As before, we must show that the loss in district θ^* outweighs any gain in district $\theta = \omega + s^*$ for ω in the support of F^p. The probability that party 1 wins district θ^* is $L\big(x(\theta^*) + \beta(x(\theta^*) - y_2'')\big)$, and the probability that party 1 wins district θ is $L\big(x(\theta) - \beta(x(\theta^*) - y_2'')\big)$. Since $x(\theta^*) > 0 > x(\theta)$ and $x(\theta^*) < -x(\theta)$ it follows from the properties of L that

$$L\big(x(\theta^*) + \beta(x(\theta^*) - y_2'')\big) - L(x(\theta^*)) > L(x(\theta)) - L\big(x(\theta) - \beta(x(\theta^*) - y_2'')\big)$$

and hence $y_2'' \neq x(\theta^*)$ increases the seat share of party 1 as desired.

Next, we show that the critical state is s along the equilibrium path. To see this, first note that party 2 can always match party 1's policy choice. Since the redistricting plan is optimal for the case in which both parties choose the same policy, it follows that the critical state can be no lower than s. Conversely, if party 1 chooses F^p and y_1, then, by the claim above, party 2's unique best response is to choose y_2. Therefore, the critical state must be equal to s.

By Theorem 1 and the fact that party 2 can match party 1's policy, this implies that in any equilibrium the redistricting plan must be equal to F^p. By the claim above, it then follows that party 1 must choose y_1 and party 2 must choose y_2 along the equilibrium path.

To establish the existence of an equilibrium in every subgame, note that for any redistricting plan $F \in \mathcal{F}$ and any policy choice for party 1, there exists an optimal policy for party 2. This follows from the continuity of party 2's payoff function.

PROOF OF COROLLARY 2:

Let p^n and s^n be the equilibrium redistricting plan and the critical state respectively. It suffices to show that $\lim_{n\to\infty} [\omega^*(p^n) + s^n] = ½$. To prove this, let $\pi^n(\theta) = L^n(x(\theta))$ and let $\phi^n(\theta)$ be as defined previously. By Lemma ϕ, the functions ϕ^n are decreasing and $\phi^n(\theta) \geq ½$ for all n, θ. Since $\lim_{n\to\infty} \pi^n(z) = 1$ for all $z > v$, $\lim_{n\to\infty} \phi^n(\theta) = ½$. Therefore, the sequence of functions ϕ^n converge uniformly to ½. Since $\omega^*(p^n) + s^n = \phi^n(\omega)$ for some $\omega \leq ½$, it follows that $\lim_{n\to\infty} [\omega^*(p^n) + s^n] = ½$ as required.

REFERENCES

Besley, Timothy, and Ian Preston. 2007. "Electoral Bias and Policy Choice: Theory and Evidence." Institute for Fiscal Studies Working Paper W07/06.

Cameron, Charles, David Epstein, and Sharyn O'Halloran. 1996. "Do Majority-Minority Districts Maximize Substantive Black Representation in Congress?" *American Political Science Review*, 90(4): 794–812.

Coate, Stephen, and Brian Knight. 2007. "Socially Optimal Districting: A Theoretical and Empirical Exploration." *Quarterly Journal of Economics*, 112(4): 1409–71.
Cox, Gary W., and Jonathan N. Katz. 1999. "The Reapportionment Revolution and Bias in U.S. Congressional Elections." *American Journal of Political Science*, 43(3): 812–40.
Cox, Gary W., and Jonathan N. Katz. 2002. *Elbridge Gerry's Salamander: The Electoral Consequences of the Reapportionment Revolution.* Cambridge: Cambridge University Press.
Diamond, Peter A., and Joseph E. Stiglitz. 1974. "Increases in Risk and in Risk Aversion." *Journal of Economic Theory*, 8(3): 337–60.
Epstein, David, and Sharyn O'Halloran. 1999. "Measuring the Electoral and Policy Impact of Majority-Minority Voting Districts." *American Journal of Political Science*, 43(2): 367–95.
Friedman, John N., and Richard T. Holden. 2008a. "Optimal Gerrymandering in a Competitive Environment." Unpublished.
Friedman, John N., and Richard T. Holden. 2008b. "Optimal Gerrymandering: Sometimes Pack, but Never Crack." *American Economic Review*, 98(1): 113–44.
Gelman, Andrew, and Gary King. 1990. "Estimating the Electoral Consequences of Legislative Redistricting." *Journal of the American Statistical Association*, 85(410): 274–82.
Gelman, Andrew, and Gary King. 1994. "Enhancing Democracy through Legislative Redistricting." *American Political Science Review*, 88(3): 541–59.
Gilligan, Thomas W., and John G. Matsusaka. 2006. "Public Choice Principles of Redistricting." *Public Choice*, 129(3–4): 381–98.
Grigg, Delia, and Jonathan N. Katz. 2005. "The Impact of Majority-Minority Districts on Congressional Elections." Paper presented at the annual meeting of the The Midwest Political Science Association, Chicago, IL.
Herron, Michael C., and Alan E. Wiseman. 2008. "Gerrymanders and Theories of Law Making: A Study of Legislative Redistricting in Illinois." *Journal of Politics,* 70(1): 151–67.
King, Gary, and Robert X. Browning. 1987. "Democratic Representation and Partisan Bias in Congressional Elections." *American Political Science Review,* 81(4): 1251–73.
Owen, Guillermo, and Bernard Grofman. 1988. "Optimal Partisan Gerrymandering." *Political Geography Quarterly,* 7(1): 5–22.
Rodden, Jonathan. 2007. "Political Geography and Electoral Rules: Why Single-Member Districts are Bad for the Left." Unpublished.
Sherstyuk, Katerina. 1998. "How to Gerrymander: A Formal Analysis." Public Choice, 95(1–2): 27–49.
Shotts, Kenneth. 2001. "The Effect of Majority-Minority Mandates on Partisan Gerrymandering." *American Journal of Political Science*, 45(1): 120–35.
Shotts, Kenneth. 2002. "Gerrymandering, Legislative Composition, and National Policy Outcomes." *American Journal of Political Science*, 46(2): 398–414.

This article has been cited by:

1. Clemens Puppe, Attila Tasnádi. 2014. Axiomatic districting. *Social Choice and Welfare* . [CrossRef]
2. Emanuele Bracco. 2013. Optimal districting with endogenous party platforms. *Journal of Public Economics* **104**, 1-13. [CrossRef]
3. Francesco De Sinopoli, Leo Ferraris, Giovanna Iannantuoni. 2013. Electing a parliament. *Social Choice and Welfare* **40**:3, 715-737. [CrossRef]
4. TODD MAKSE. 2012. Strategic Constituency Manipulation in State Legislative Redistricting. *Legislative Studies Quarterly* **37**:2, 225-250. [CrossRef]
5. KEISUKE NAKAO. 2011. RACIAL REDISTRICTING FOR MINORITY REPRESENTATION WITHOUT PARTISAN BIAS: A THEORETICAL APPROACH. *Economics & Politics* **23**:1, 132-151. [CrossRef]