Optimal Gerrymandering

John N. Friedman and Richard T. Holden*

September 30, 2005

Abstract

Standard intuitions for optimal gerrymandering involve concentrating one's extreme opponents in "unwinable" districts ("throwing-away") and spreading one's supporters evenly over "winnable" districts ("smoothing"). These intuitions are not robust and depend crucially on arbitrary modelling assumptions. We characterize the solution to a problem in which a gerrymanderer observes a noisy signal of voter preferences from a continuous distribution and creates N districts of equal size to maximize the expected number of districts which she wins. We show that "throwing-away" districts is not generally optimal, nor is "smoothing." The optimal solution involves creating a district which matches extreme "Republicans" with extreme "Democrats," then continuing to match toward the center of the signal distribution. We show that the value to being the gerrymanderer is increasing in the extremity of voter preferences, the quality of the signal, and the number of districts.

Keywords: Endogenous political institutions, gerrymandering, redistricting.

JEL Codes: D72, H10, K00.

---

*Department of Economics, Harvard University. Correspondence: Richard Holden, Littauer Center, 1875 Cambridge Street, Cambridge, MA, 02138. email: rholden@fas.harvard.edu; jnfriedm@fas.harvard.edu. We would like to thank Philippe Aghion, Alberto Alesina, Rosalind Dixon, Benjamin Friedman, Noah Feldman, Drew Fudenberg, Edward Glaeser, Christine Jolls, David Laibson, Jesse Shapiro, Andrei Shleifer and Jeremy Stein for helpful suggestions, and also participants in seminars at Harvard Law School and Harvard University. Friedman acknowledges financial support from the National Science Foundation.

1 Introduction

One of the more curious features of American democracy is that electoral boundaries are drawn by political parties themselves. In order to ensure a notion of equal representation, the Constitution of the United States provides that "Representatives and direct Taxes shall be apportioned among the several States which may be included within this Union, according to their respective Numbers."[1] Since populations change over time, the Constitution also provides a time frame according to which representation shall be adjusted - "...within every subsequent Term of ten Years, in such Manner as they shall by Law direct"- where the "they" are the several states. In practice, this leaves the process of redistricting to state legislatures and governors.

History has shown that political parties act in their own interests; redistricting is no exception, and the advantages gained from it are large. From Massachusetts's Elbridge Gerry in 1812 (after whom the term "Gerrymander" was coined), to the recent actions of Texas's Tom DeLay, American politicians have long used the redistricting process to achieve partisan political ends. Most recently, the much publicized Republican redistricting in Texas in 2003 caused four Democratic Congressman to lose their seats and would have been even more extreme but for the Voting Rights Act, which effectively protected nine Democratic incumbents. Other particularly stark current examples include Florida, Michigan and Pennsylvania - states which are evenly divided, but whose Congressional delegations collectively comprise 39 Republicans and 20 Democrats. Democrats are also familiar with the practice; though President Bush won Arkansas by more than 10 points last November, the state's Congressional delegation, bolstered by the Democratic state legislature's redistricting in 2001, contains three Democrats and one Republican.

Although gerrymandering using unequal district sizes or racial characteristics is unlawful, partisan gerrymandering remains legal, though controversial. In *Davis v. Bandemer* (1986), the Supreme Court declared partisan gerrymandering inimical to norms of fair and equal representation; but the majority was unable enunciate a workable test for where redistricting stops and gerrymandering begins. Nearly two decades later, despite numerous attempts to find such a standard, four members of the court (Chief Justice Rehnquist and Justices O'Connor, Scalia and Thomas)

[1] Article I, section 2, clause 3.

found in *Vieth v. Jubelirer* (2004) (a 4-1-4 decision) that the test laid down in *Bandemer* was not practicable, in that it gave no guidance to legislatures and lower courts, and, absent such a test, partisan redistricting was not justiciable[2].

In the wake of this decision and the controversial Texas redistricting in 2003, there has been renewed interest in legislative reform to change the partisan nature of redistricting. Currently, two states, Iowa (since 1980) and Arizona (since 2000), include non-partisan commissions in their decennial redistricting processes, but only Arizona completely excludes political bodies. More than twenty states have considered similar amendments in the past decade, though, and movements advocating such changes seem to be gaining momentum. Bruce E. Cain, director of the Institute of Governmental Studies at University of California at Berkeley, recently commented that:

> You cannot believe the number of people and organizations across the country that are focusing on this redistricting issue... It seems like it's poised to become, for the reform community, the equivalent of McCain-Feingold[3].

Most recently, Governor Arnold Schwarzenegger of California proposed that retired judges take charge of the redistricting process. But despite the great impact of gerrymandering on the American political system and the surge of recent interest in reform, few authors have attempted to understand the basic incentives at work.

In this paper, we view the issue of redistricting through the lens of an economist concerned with the endogenous formation of political institutions. In particular, we frame the issue as a maximization problem by the gerrymander where the choice variables are the allocations of voters to districts. In contrast, most previous analyses model the problem as a trade-off between "biasedness" - the degree to which an evenly divided population would elect an uneven slate of legislators - and "sensitivity" - the responsiveness of the share of seats held by a party to the share of voters supportive of that party (Owen and Grofman, 1988; Sherstyuk, 1998; Cox and Katz, 2002). In these models, the gerrymanderer optimally concentrates those least likely to vote for her in districts which are "thrown away", and spreads remaining voters

[2] "...the legacy of the plurality's test is one long record of puzzlement and consternation.", Scalia J.

[3] Adam Nagourney, "States See Growing Campaign to Change Redistricting Laws," *New York Times*, February 7, 2005.

evenly over the other districts, which are "smoothed." A major limitation of these models is that they are not micro-founded; the gerrymanderer chooses properties of the redistricting plan, as a whole, rather than the placement of voters into districts. Since there is no one-to-one mapping from these aggregate characteristics to individual district profiles, there is no guarantee that the solution from these models is actually optimal.

Gilligan and Matsusaka (1999) take an alternative approach, instead analyzing a micro-founded model in which individuals with known party affiliations vote for those parties with probability one. Since one party wins a district comprising $n+1$ of its supporters and n opponents with certainty, the optimal strategy is to make as many districts like this as possible. Indeed, if one party holds bare majority of the population, then they win all districts! Though the assumptions of observability and deterministic voting simplify the analysis greatly, they clearly do so at the expense of realism.

Shotts (2002) considers the impact of majority-minority districting. He develops a model with a continnum of voters and imposes a constraint he calls the "minimum density constraint." This requires the gerrymanderer to put a positive measure of all voter types in each distict. This appears to be a reduced form way of modelling the fact that the gerrymanderer observes a noisy signal of voter preferences. We explicitly model the signal, and are thus more general. We note that the minimum density constraint *a priori* rules out what we show is, in fact, the optimal strategy - matching slices.

We analyze a model in which there is a continuum of voter preferences, and where the gerrymanderer observes a noisy signal of these preferences. We show that the optimal strategy involves creating districts by matching increasingly extreme blocks of voters from opposite tails of the signal distribution. This finding contrasts with the bulk of the previous literature in two ways: First, we show that extreme Democrats should be matched in a district with extreme Republicans, rather than concentrated in a district which Republicans concede to be unwinnable and "throw away." Intuitively, extreme Democrats can be best neutralized by matching them with a slightly larger mass of extreme Republicans. Second, we show that it is better to put extreme Republicans and moderate Republicans in separate districts. This contrasts with the "smoothing" intuition, which calls for the creation of identical profiles among districts which the gerrymanderer expects to win. Intuitively, since district composition determines the median voter, smoothing districts makes inefficient use of extreme Republicans as right-of-the-median voters in many districts, rather than having them be the median in some districts.

This analysis is a first step toward a more complete understanding of the phenomenon of gerrymandering. There are important issues which this paper does not address: most notably, we focus exclusively on partisan incentives, to the exclusion of the motivations of incumbents (i.e. incumbent gerrymandering). We also abstract from geographical considerations, such as the legal requirement of contiguity (which, we will argue, may be surprisingly unimportant), as well a preference for compactness or the recognition of communities of interest. Finally, the gerrymanderer in our model does not account for the potential limitations of "race conscious" redistricting. Of course, this does not mean that racial and partisan gerrymandering are distinct phenomena. Given that race is a component of the signal of voter preference observed by the gerrymanderer, there may be circumstances where they are essentially the same practice. Ultimately this is an empirical question, which depends on the distribution of voter preferences across voter characteristics. These issues are further explored in Section 7.

The remainder of the paper is organized as follows: Section 2 details the legal and institutional backdrop against which redistricting takes place. In Section 3 we walk through some basic examples which illustrate the primary intuitions of the solution to our more general model, which we present in Section 4 along with comparative statics. Section 5 reports the result of a number of numerical examples of the model in order to illustrate further the optimal strategy and its comparative statics. In Section 6 we sketch a number of extensions to the basic model, including uncertain voter turnout, incumbent advantages, and alternative partisan objective functions. Finally, Section 7 contains some concluding remarks and suggests directions for future research.

2 Background: Legal Decisions and Political Realities[4]

Though the process of redistricting was politicized in America as early as 1740 (in favor of the Quaker minority in the colony of Pennsylvania), the modern history of gerrymandering begins with the landmark Supreme Court decision *Baker v. Carr* (1962). Especially in the south, states had not redrawn Congressional districts after each decennial census, as mandated by the U.S. Constitution. Since population growth was much greater in urban areas, this inertia served to dilute the urban vote - often poor and black - and enhance the political power of rural white voters who traditionally supported the Democratic Party. After the 1960 census, the population disparities between Congressional districts had become as great as 3 to 1 in Georgia (and as extreme as 1000 to 1 for state legislature seats in some states). The decision in *Baker* declared that challenges to such districting plans were justiciable, and two years later the Court clarified its position on the standard for unlawful redistricting plans, stating in *Wesberry v. Sanders* that only Congressional districts with populations "as nearly equal as possible" were acceptable under the Equal Protection clause.[5] Furthermore, Federal District courts were empowered, as part of their remedial discretion, to draw district boundaries themselves should a state prove either unable or unwilling to produce a satisfactory plan.

Consensus over the practical implications of the Court's decisions solidified over the next 15 years. Though Federal District courts initially experimented with strict upper bounds on the maximum population deviation across districts, by the late 1970s states were subject to a more flexible set of criteria, in which concerns such as the compactness of districts or the preservation of "communities of interest" justified small deviations in representation. As of 1980, though, population equality across districts was the only constraint on redistricting. In fact, the 1980 round of redistricting demonstrated that strict adherence to population equality gave states great latitude to skew election results by other means. For instance, the California redistricting plan, despite blatantly favoring Democrats, withstood court challenge due to near population equality, while courts rejected the less partisan New Jersey

[4] This section details the legal and political backdrop against which gerrymandering occurs today. Readers uninterested in or already familiar with this material may wish to skip directly to the analysis in Section 3.

[5] See *Westberry v. Sanders* 376 US 1 (1964). The court applied a similar standard to districts for statewide legislative bodies in *Reynolds v. Sims* 377 US 533 (1964) and for general purpose local governments in *Avery v. Midland County* 390 US 474 (1968).

plan due to its unacceptable maximum population deviation of 0.6984%.[6] Perhaps spurred on by the outcomes of the 1980 round of redistricting, the Court attempted to expand the list of justiciable claims for redistricting plans to include both racial and partisan motivations. These efforts have proved less lasting than those based on population equality.

In the 1990's, debates around gerrymandering shifting to the issue of "race conscious" redistricting. While it had long been clear that intentional dilution of the voting strength of racial minorities violated the Equal Protection clause, it was less clear that states could draw boundaries such that racial minorities could elect their preferred candidates (Issacharoff, Karlan and Pildes, 2002). In a number of cases, culminating in *Shaw v. Reno* (1993), the Court found that redistricting plans would be held to the same strict scrutiny with respect to race as other state actions. In practice, this means that, once plaintiffs demonstrate that racial concerns were a "predominant factor" in the design of a districting plan, the plan is illegal unless the state can justify the use of race and show that such factors were considered only when necessary. This places a heavy burden on the states. Some Federal courts initially interpreted these decisions as requiring states to ensure minority representation through the creation of minority-majority districts, but the Supreme Court declared that this practice would violate Section 2 of the Voting Rights Act. In more recent cases, the Court continues to downplay the importance of racial considerations; for instance, litigation surrounding the 1991 North Carolina redistricting ended when the Court ruled, in *Easley v. Cromartie* (2001) that partisan concerns, not racial concerns, "predominated" in the construction of the heavily black and democratic 12th district, and thus the plan was legal.

The history of attempts to ban partisan gerrymandering have proved less successful still. In *Davis v. Bandemer*, the Supreme Court attempted to limit the impact of partisan concerns in redistricting processes by stating that such claims were, in theory, justiciable (though they did not decide the merits). Though the years following this decision saw many attempts to define the level and shape of such a standard, there was little agreement, and no claim of partisan gerrymandering ever succeeded. In *Vieth v. Jubelirer*, four member of the Court found that such

[6] See *Karcher v. Daggett* (1983). In gerrymandering jargon, the "population deviation" of each district is the difference between the actual district population and the ideal district population, expressed in terms of percent of the ideal population. The "maximum population deviation" is the difference between the greatest and lowest population deviations.

attempts were doomed. While *Bandemer* is still good law, the future justiciability of partisan gerrymandering claims seems far from assured.

The current reality of political redistricting reflects the past forty years of case history. States now use increasingly powerful computers to aid in the creation of districts, and, accordingly, *Baker*'s "as nearly equal as possible" population requirement is extremely strict. A Pennsylvania redistricting plan was struck down in 2002 for having one district with 19 more people than another without justification! On the other hand, the law does allow for some slight deviations, provided there is adequate justification. In Iowa, for instance, Congressional districts must comprise whole counties; the current maximum population deviation of the Iowa redistricting plan is 131 people, but the legislature rejected an earlier plan with a 483-person deviation. Such cases are not common, though. The current Texas districting plan is more representative and has, to integer rounding, exactly equal population in each district.

In addition to the equal population requirement, districts must be contiguous. This requirement first appears in the Apportionment Act of 1840, though it was standard long before then. While technology has tightened the population equality constraint, computers have effectively loosened the contiguity requirement, as legislators can now draw districts more finely than ever before. In the 1970s, districting plans were laborious to create and difficult to change, as each required hours of drawing on large floor-maps using dry-erase markers; now lawmakers use Census TIGERLine files to create and analyze many alternative districting schemes both quickly and accurately. Contiguity has been stretched to the limit in such recent cases. Florida's 19th, 22nd, and 23rd districts, shown in Figure 1 (in the Appendix), are one such case. The 22nd comprises a coastal strip not more than several hundred meters wide in some places but ninety miles long, while tentacles from the 22nd and 23rd intertwine to divide the voters of West Palm Beach and Fort Lauderdale. Even more striking is the shape of the Illinois 4th (shown in Figure 2), drawn to include large Hispanic neighborhoods in the North and South of Chicago but not much in between. Each of these districts is, in some places, no more than one city block wide, and such necks are often narrower than 50 meters.

Digital tools have also aided the rise of a partisan practice known as "kidnapping," in which the home of an incumbent of one party is included in a different (and typically far-off) district using a long narrow connecting strip. This maneuver deprives the candidate of much of her incumbency advantage and may pit her against another incumbent to lower further her chances of reëlection. Four democratic Congressmen from Texas lost their seats through such tactics in the 2004. Because of the obvious spacial aspect of kidnapping, we do not consider the practice in our analysis.

Each state has its own procedure for redrawing district boundaries. In most states, redistricting plans are standard laws, proposed by the members of the legislature and subject to approval by the legislatures and the governor. Arizona and Iowa delegate redistricting to independent commissions, though in Iowa legislators must still approve the plan and may edit proposed schemes after several have been rejected. In 2001, for instance, the legislature rejected the first proposed plan along partisan lines because Republicans thought the plan was not favorable enough.[7] Arizona and Iowa also instruct their redistricting commissions to make districts compact, respect the boundaries of existing "communities of interest," and use geographic features and existing political boundaries to delineate districts "to the extent practicable." Finally, Arizona also mandates that "competitive districts should be favored where to do so would create no significant detriment" to other objectives.[8] No other states have explicitly defined redistricting goals along these lines.

There are three key messages to understand from the backdrop against which gerrymandering takes place. First, contiguity may well not be a binding constraint because of the fine lines which gerrymanderers use to create districts. Second, other spatial/geographic concerns such as compactness and communities of interest have found little legal traction. As such, they are really not constraints on gerrymanderers. Third, the Court has consistently considered partisan and racial gerrymandering to be analytically distinct - *Cromartie* even going so far as to allow racial gerrymandering if it is not deemed the predominant motive. The first two of these points suggest that spatial/geographical considerations are not first-order concerns. Accordingly, our model omits them. The third rests on the premise that signals of voting propensity and race are sufficiently uncorrelated that an optimal gerrymandering strategy does not conflate the two issues. This is a point to which

[7] "Senate Rejects Districts," *Des Moines Telegraph Herald*, May 3, 2001.

[8] See Arizona Proposition 106, and 1981 Iowa Acts, 2nd Extraordinary Session, Ch. 1.

we return later in the paper.

3 Some Simple Examples

In order to illustrate the intuition behind the theory in this paper, we now provide some very simple examples that capture the basic features of the more general model in Section 4.

3.1 Example 1

Consider the problem faced by a gerrymanderer in a state in which a population of voters have single-peaked preferences over a one-dimensional policy space. We assume that each voter has bliss point β, and that, across the population, β is distributed uniformly on $[-1, 1]$. These assumptions imply that, in a two party election, each voter supports the candidate located closest to her on the ideological spectrum. To begin, we assume that the gerrymanderer can directly observe β for each voter. We assume that all candidates - the right-wing "Republican" candidate and the left-wing "Democrats" - locate at 0, and so the percent of votes captured by the Republican candidate in any election is simply the proportion of voters to the right of 0.

The gerrymanderer - suppose she is a Republican - must break up the population into equal-sized districts in which different elections take place with the goal of maximizing the expected number of seats won by her party. Since we abstract from geographic concerns here, the gerrymanderer can match together any pieces of the population into a district. Suppose, for simplicity, that the gerrymanderer must form 2 districts, so that each district must comprise mass $\frac{1}{2}$ of voters. Since all voters for whom $\beta \geq 0$ support the Republican candidate with certainty, Republicans win all districts containing $\frac{1}{4}$ or more mass of such voters.[9] From Gilligan and Matsusaka (1999), the optimal gerrymander makes exactly half of the voters in each district have preferences $\beta \geq 0$; in this basic setup, Republicans win each district with certainty. It does not matter which right-wing voters go into each district.

[9] For the sake of simplicity, we resolve all "ties" in this example in favor of the Republican candidate. Voters with $\beta = 0$ support the right-wing candidate, and if the candidates have equal vote shares the Republican wins.

3.2 Example 2

We now add some noise to the preferences in example 1. Suppose that, after candidates are positioned at 0, an aggregate preference shock A affects the population so that preferences are now single-peaked about $\hat{\beta} = \beta - A$. The gerrymanderer observes only β and not A or $\hat{\beta}$. Suppose that A is distributed uniformly on $[-1, 1]$. While voters for whom $\beta > 0$ now vote for the right-wing candidate in expectation, only those for whom $\beta = 1$ support the Republican candidate with certainty; a voter with $\beta = 0.5$, for instance, only prefers the right-wing candidate if $A < 0.5$, which happens 75% of the time.

In this example we can make a sharper prediction about the form of the optimal gerrymander. Half of the voters in each district should have $\beta > 0$, but it now matters which of these voters go into which district. The optimal gerrymander groups all extreme voters for whom $\beta \in [0.5, 1]$ into one district (denoted as District 1) and more moderate right-wingers with $\beta \in [0, 0.5]$ into District 2. These blocks of right-wing voters are then grouped with any mass of voters for whom $\beta < 0$; since the preference of the median voter in each district ($\mu_1 = 0.5$ in District 1 and $\mu_2 = 0$ in District 2) is already determined, the composition of the left-wing voters does not matter. Republican candidates now win District 1 with probability 0.75 and District 2 with probability 0.5. Any other distribution of right-wing voters between the two districts (with $\frac{1}{4}$ mass to each) would dilute the power of the extreme right-wing voters by wasting some in District 2, since that median voter would still have $\beta = 0$ while the preferences of the median voter of District 1 would fall. Only by concentrating the most extreme right-wing voters together can the gerrymanderer make the most effective use of her supporters.

3.3 Example 3

Finally, suppose (in addition to the setup in the second example) that individual preferences are measured with noise by the political parties. That is, let the gerrymanderer only observe s, a signal of preferences, instead of β itself. Across the population, let s be distributed uniformly on $[-1, 1]$, and let $\beta \mid s$ be distributed uniformly on $[s - 0.5, s + 0.5]$, with an independent draw of β for each voter with a given signal s. Suppose the gerrymanderer creates districts as above (grouping voters for whom $s \in [0.5, 1]$ into District 1 and $s \in [0, 0.5]$ into District 2), and, furthermore, groups the most extreme left-wing voters into District 1 and the others in District 2. Because the measurement of preferences is noisy, the median voter in District 1 falls to $\mu_1 = 0$; the Republicans gain no advantage over proportional representation. Intuitively, the Republicans are "cutting it too close" in District 1. Although District 1 contains the most extreme right-wing voters, there are only $\frac{1}{4}$ mass of them, and so the most left-wing voter with a right-wing signal is the median voter. Since some of those right-wing voters end up with more moderate preferences than their signal suggested, the median voter in the district is a moderate.

Instead, consider a gerrymander which groups all voters with $s \in [p, 1]$ into District 1 and $s \in [0, p]$ into District 2. Because of the intuition developed in the second example, this districting scheme still keeps the most extreme right-wing voters together. Now, though, the Republicans have more than just a bare majority of supporters in District 1, reducing the problem caused by preference mis-measurement above.

To complete this optimal districting, the gerrymanderer must allocate the left-wing voters. Her problem here is exactly opposite that faced with the right-wing voters: she must decide how best to *neutralize* the voting power of the extreme left-wingers. The key to this problem is that, since the majority of District 2 voters are left-wingers (assuming $p > \frac{1}{2}$), μ_2 is far more sensitive to the allocation of the these voters than μ_1. Thus, the optimal gerrymanderer should concentrate those least likely to vote for the Republican candidate into District 1, where they do not affect the median voter.

Combining these insights, consider a districting plan such that voters for whom $s \in [-1, -1 + p] \cup [p, 1]$ make up District 1 and the rest are placed in District 2. The particular distributional assumptions made above imply that

$$\mu_1 = p + \sqrt{1 - 2p} - \frac{1}{2} \qquad \text{and} \qquad \mu_2 = p - \frac{1}{2}.$$

The optimal gerrymander sets $p^* = \frac{3}{8}$; Republican candidates win 1.125 seats in expectation. By including more right-wingers in District 1, μ_1 becomes less sensitive to the mis-measurement of preferences, and thus increases quite a bit, while μ_2, which depended less on the precision of the signal, does not decrease by as much.

Furthermore, the right-wing voters of District 1 determine that $\mu_1 = \sqrt{\frac{1}{4}} - \frac{1}{4}$, and so the inclusion of the most extreme left-wingers has no effect. If, for instance, the gerrymanderer had included these least favorable voters into District 2 and placed voters with $s \in [-1 + p, -1 + 2p]$ into District 1, μ_2 would fall while μ_1 would not change.

These three simple examples illustrate how key features of an optimal partisan gerrymander differ from the standard "throwing away" and "smoothing" intuitions. For a Republican gerrymanderer, it is not best to "smooth" extreme and moderate right-wing voters across many districts; rather, one should concentrate the most extreme right-wingers into a single district in order to not waste them all as right-of-median voters. Second, it is not efficient to concentrate those least likely to vote for one's candidate into a district that is "thrown away"; instead, these extreme left-wingers voters are best countered by matching them with a greater number of extreme right-wingers.

We now turn to our model, which provides a more general characterization of the optimal partisan gerrymander, but the intuitions brought out in our examples are still prominent. Indeed, under certain regularity conditions, the optimal districting scheme has exactly the same form as in the final example above, matching increasingly extreme slices of voters from opposite sides of the signal distribution for the population.

4 The General Model

4.1 Overview

There are two parties, D and R, which can be interpreted as the Democratic Party and the Republican Party. One of these parties (for simplicity we assume it to be R) is the gerrymanderer and creates district profiles. There is a continuum of voters, each of whom is characterized by a policy preference parameter. The gerrymanderer does not observe this parameter but, instead, receives a noisy signal of it. Also, she observes the posterior distribution of policy preference parameters conditional on the signal. We will sometimes refer to the marginal distribution of the signal as the "signal distribution". Thus, her problem is to create N voting districts by allocating voters from the signal distribution. Her objective is to maximize the expected number of districts won. We determine the probability that each party wins a district by the median voter in that district. The only constraints we place on the gerrymanderer are that: (i) each voter must be allocated to one and only one district; and (ii) all districts must contain an identical mass of voters.

4.2 Statement of the Problem

There is a continuum of voters who differ in their political preference. We assume that each voter has single-peaked preferences about a bliss point $\beta \in \mathbb{R}$ which varies across the population. These preferences are not observed by the gerrymanderer, who instead receives a noisy signal, $s \in \mathbb{R}$. Let the joint distribution of β and s be given by $F(\beta, s)$ on support $\mathbb{R}^2$. Let player R be the gerrymanderer. Let R have a Bayesian posterior $G(\beta \mid s)$ for the distribution of preferences given an observed signal. We refer to this distribution as the "conditional preference" distribution. We assume that both F and G are absolutely continuous. Define the marginal distribution of s as:

$$H(s) = \int F(\beta, s)\, d\beta$$

Since there are a continuum of voters we can interpret H not only as characterizing a single draw from the population of voters, but also the mass of voters in the population. We refer to H as the "signal distribution". R allocates mass from this distribution in order to form districts. Let H have (continuous) probability mass function $h(s)$. Without loss of generality, let the median of s in the population equal 0.

Since preferences are single-peaked, voters choose the candidate closest to them on the ideological spectrum. Furthermore, we assume that each candidate positions herself at the median of the signal distribution for the population, which is 0 by assumption. Thus, all voters to the right support the Republican candidate, while those to the left support the Democrat. As a reduced form representation of electoral uncertainty, we assume that, in each election *after* the candidates are positioned, there is an aggregate shock decreasing all preferences by A. Thus, if the median voter in district n has preference μ_n, she votes for the Republican candidate if and only if $A \leq \mu_n$, which occurs with probability $B(\mu_n)$, where $B(\cdot)$ denotes the c.d.f of A. We assume that A can take any value in $\mathbb{R}$ with positive probability, so that B is strictly increasing. One can also think of A as an "electoral breakpoint" such that voters positioned above (to the right) of the realization of the breakpoint vote for the Republican candidate, while those on the left vote democratic. Importantly, once the breakpoint is determined, all uncertainty is resolved and the position of voters relative to A determines for whom they vote with certainty. The uncertainty about whom a particular voter will vote for comes from the fact that A is stochastic.

Since we assume that all candidates locate at the population median, we necessarily imply that there is nothing "local" about an election. Though perhaps counter-intuitive, research suggests that this may not be far from the truth. Ansolabehere, Snyder and Stewart (2001) argue that, while district-to-district competition may exert some influence on the candidate platforms, the effect is "minor compared to the weight of the national parties." Allowing for state-to-state differences would surely leave even less variation in local platforms. Similarly, Lee, Moretti and Butler (2004) demonstrate that exogenous shifts in electoral preferences do not affect the menu of candidates offered to voters, perhaps because politicians have no way to credibly commit to campaign promises. We discuss the effects of certain departures from this assumption in Section 6.

R divides the population into N equal-sized districts to maximize the expected number of seats won in the election. Let $\psi_n(s)$ denote the mass of voters from the population placed in district n. Formally, R solves the program

$$\begin{aligned} &\max_{\{\psi_n(s)\}_{n=1}^N} \left\{ \frac{1}{N} \sum_{n=1}^{N} B(\mu_n) \right\} \\ s.t. \quad &\int_{-\infty}^{\infty} \psi_n(s)\, ds = \frac{1}{N} \qquad, \forall n, s \\ &\sum_{n=1}^{N} \psi_n(s) = h(s) \\ &0 \leq \psi_n(s) \leq h(s) \quad, \forall n, s \end{aligned} \tag{1}$$

where

$$\mu_n = \hat{\beta} \text{ s.t. } \int_{-\infty}^{\infty} G\left(\hat{\beta} \mid s\right) \psi_n(s)\, ds \equiv \Gamma_n\left(\hat{\beta}\right) = \frac{1}{2N}. \tag{2}$$

It will be useful to define the following for notational purposes:

$$\gamma_n(\beta) = \frac{\partial \Gamma_n(\beta)}{\partial \beta}. \tag{3}$$

Note that, given a district profile $\psi_n(s)$, equation (2) determines μ_n with certainty. Though R could not identify any single voter as the median voter in a district, there is nothing stochastic about the preference parameter of the median voter.[10]

In order to analyze the problem it is necessary to place some structure on the conditional distribution of preferences. The first restriction we require is that the signal be informative in the following sense.

Condition 1 (Informative Signal Property) *Let* $\frac{\partial G(\beta|s)}{\partial s} = z(\beta \mid s)$. *Then*

$$\frac{z(\beta \mid s')}{z(\beta \mid s)} < \frac{z(\beta' \mid s')}{z(\beta' \mid s)}, \forall s' > s, \beta' > \beta$$

This property is similar to the Monotonic Likelihood Ratio Property due to Karlin and Rubin (1956) (see also Milgrom, 1981). In fact, if a higher signal simply shifts the mean of the conditional preference distribution then this property is equivalent to MLRP. When this is the case, the condition essentially states that higher and higher signals (more right-wing) are more and more likely to come from voters who have underlying preferences which are further to the right. Many common distributions satisfy it, including: the normal, exponential, uniform, chi-square, Poisson, binomial, non-central t and non-central F. If a higher signal also changes the shape of the conditional distribution then this property, like MLRP, becomes less intuitive. It does, in general, imply First Order Stochastic Dominance[11], and as such rules out cases where observing a higher signal makes *both* the probability of the voter being extreme left-wing *and* the probability of being extreme right-wing increase.

The second condition we require is unimodality.

[10] This model structure is isomorphic to the inclusion of further levels of uncertainty between signals and preferences. For instance, suppose that the gerrymanderer believed that, with 50% probability, preferences had a conditional distribution $G_1(\beta \mid s)$, and otherwise they were conditionally distributed as $G_2(\beta \mid s)$. Equation (2) would then become $\int_{-\infty}^{\infty} \frac{1}{2}[G_1(\mu_n \mid s) + G_2(\mu_n \mid s)]\psi_n(s)\,ds = \frac{1}{2N}$, which is isomorphic to our original problem if instead $G(\beta \mid s) = \frac{1}{2}[G_1(\mu_n \mid s) + G_2(\mu_n \mid s)]$.

[11] As MLRP always does.

Condition 2 (Central Unimodality) *$g(\beta \mid s)$ is a unimodal distribution where the mode lies at the median.*

Also note that, without loss of generality, we can "re-scale" s such that $s = \max_\beta g(\beta \mid s)$. Though many distributions that satisfy Condition 1 are unimodal, some are not, and we rule these out. Furthermore, Condition 1 implies that the mode of $g(\beta \mid s)$ must lie below the mode of $g(\beta \mid s')$ if $s < s'$. We can thus "re-label" the signals such that the mode of $g(\beta \mid s)$ lies at s. The two properties in Condition 2, taken together, intuitively imply that, conditional on signal s, preferences are distributed "near" s and not elsewhere.

4.3 Solution

We characterize the optimal Gerrymandering strategy in four steps. These steps describe how to "chop-up" the H distribution optimally. This is a complicated problem due to the size of the space over which the optimization takes place. Lemma 1, below, shows that we can, without loss of generality, dramatically simplify this problem by restricting attention to a much smaller space. (The Appendix contains all proofs).

4.3.1 Step 1: Slices and Parfaits

Lemma 1 *Suppose Condition 1 holds. Then voters with the same signal can appear in two different districts only if the median voter in those two districts is identical. That is, for $n \neq m$ and $s \neq s'$, $\psi_n(s) > 0$ and $\psi_m(s) > 0$ and $\psi_n(s') > 0$ and $\psi_m(s) > 0 \Longrightarrow \mu_n = \mu_m$.*

The first step demonstrates that one can restrict the search for the optimal gerrymander from the space of all functions to just vertical slices and "parfaits" of the signal distribution h. Furthermore, parfaits must themselves combine to form vertical slices, and so the vertical slices are the primary building blocks for the optimal strategy.[12] Figure 3, below, shows an example of a districting scheme.

[12] These vertical slices can actually be split many different ways between the two districts so that both the median μ and the density of district preferences $\gamma(\beta)$ at the median remain the same. One such way to split these vertical slices between the two districts, perhaps the simplest way, is a "parfait," an equal split of $h(s)$ for all s in the districts. Since all such splits are equivalent in their implications for the value function, we can, without loss of generality, consider only this simplest split.

Districts 1 through 3 are each formed by the union of an upper and a lower slice, while districts 4 and 5 are "parfaits." Formally, a parfait refers to allocating mass from H to a number of districts such that those districts have the same median. (In step three we will show that, in fact, parfaits are not optimal, therefore leaving slices as the optimal strategy.)

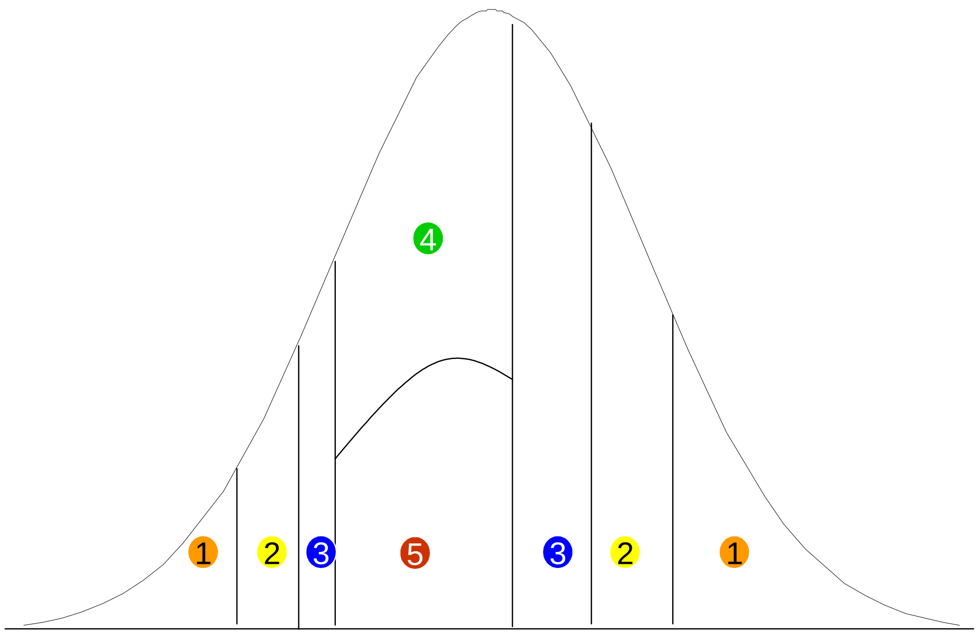

Figure 3: Slices and Parfaits

The intuition behind this result is very similar to that in our second example in Section 3 above. Since the gerrymanderer cares only about the medians of the districts, it is inefficient to spread voters with the same signal over multiple districts. In that case, extreme right-wing voters become right-of-median voters in many districts, while the actual medians remain low. Instead, it is better to concentrate voters of similar signals together in districts so as to maximize the power of their votes.

4.3.2 Step 2: The "Monotonicity" of Slices

Lemma 2 *Suppose Condition 1 holds and consider two districts, m and n, such that $\mu_n < \mu_m$. Consider any two voter types, $s_1', s_2' \in \psi_n$, in district n. Then district m can include no intermediate voter type $s \in [s_1', s_2']$, so that $s \notin \psi_m$.*

Lemma 2 shows that, in the optimal gerrymander, the voters in higher-median districts must lie outside - that is, have more extreme preferences - those in lower-median districts. The intuitions here are very similar to those discussed above. First, the same logic behind Lemma 1 - that extreme right-wing voters should be concentrated to maximize their voting strength - implies that the optimal districting scheme should place an unbroken mass of voters with higher signals into the higher-median district rather that alternate smaller slices into all districts. Second, the "extreme matching" logic from our third example in Section 3 implies that higher district medians are least sensitive to the inclusion of extreme left-wing voters in the district. Intuitively, once a district comprises more than half diehard Republicans, why not fill the rest of the district with extreme Democrats who cannot affect the electoral outcome.

4.3.3 Step 3: No Parfaits

Lemma 3 *Suppose that Conditions 1 and 2 hold. If $m \neq n$, then $\mu_m \neq \mu_n$.*

This penultimate step rules out "parfaits," as defined above. Parfaits appeared stable above because the split equated both the medians and the sensitivity of the median to changes across the two districts. But such an arrangement is a hairline case. One can reallocate mass between two such districts to maintain the equality of medians but make one district more sensitive to change than the other. Then a profitable deviation exists which lowers the less sensitive median by some but increases the other by more. Hence, parfaits cannot be optimal.

Once again, the driving intuition in this case in that of concentrating extreme voters together to maximize their electoral power. In a way, "parfaits" are the least efficient use of extreme voters, and so it cannot be surprising that they are not optimal. Thus, the optimal gerrymander must contain *only* vertical slices of the signal distribution h that do not violate the ordering restriction from Lemma 2.

4.3.4 Step 4: No Intermediate Slices

Lemma 4 *Suppose Conditions 1 holds and consider three districts m,n, and p such that $\mu_m > \mu_n > \mu_p$. Now fix $h(s)$ and N. Then for a sufficiently precise signal there does not exist a voter type $s^* \in \psi_m$ such that $s' > s_m > s''$ where $s' \in \psi_n$ and $s'' \in \psi_p$.*

This final step expands Lemma 2 by showing that voters in a higher-median district cannot lie within the set of all voters in lower-median districts. That is, by ruling out cases like that in Figure 4, it shows that optimal districts must comprise either a single slice or two slices matching mass from opposite tails of the distribution. The intuition is very similar to that of Lemma 2, that lower medians (such as those in Districts 2 and 3 in Figure 4) are more positively affected by the inclusion of moderate instead of extreme left-wing voters. On the other hand, the higher medians (such as that of District 1) are hardly lowered by the substitution of extreme left-wingers. In order for these arguments to hold, though, the signal distribution must have high enough quality. If it does not then intermediate slices are possible.

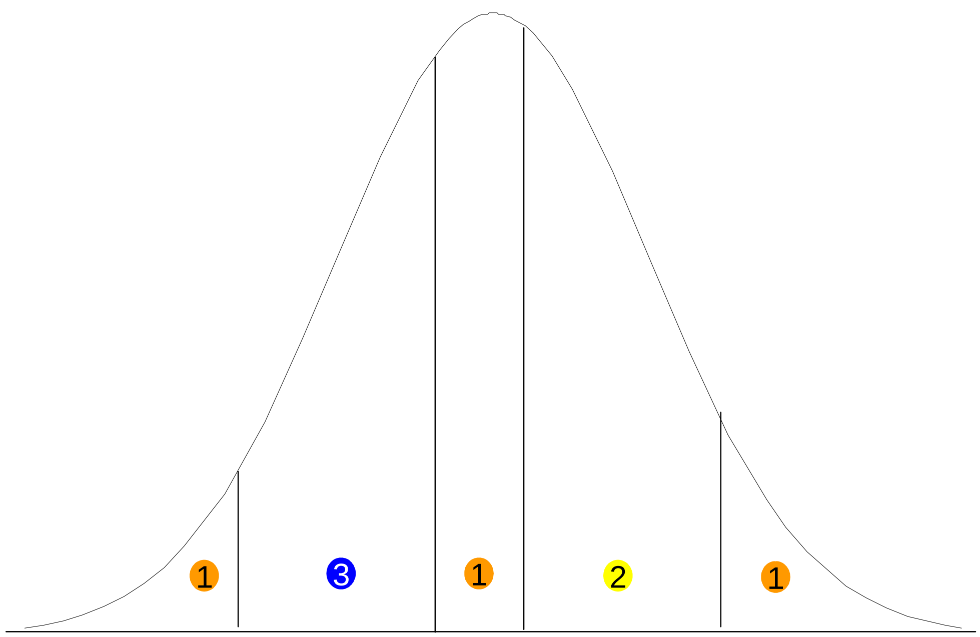

Figure 4: An example of a strategy ruled out by Lemma 4.

With these preparatory steps in place we can now provide a complete characterization of the optimal strategy for the gerrymanderer.

Proposition 1 *Suppose that Conditions 1 and 2 hold, and that the signal distribution is of sufficiently high quality (as defined in Lemma 4). Consider a districting plan with N districts labelled such that $\mu_m > \mu_n$ if $m < n$. This plan is optimal if and only it can be characterized by "breakpoints" $\{u_i\}_{i=1}^{N-1}$ and $\{l_i\}_{i=1}^{N-1}$ (ordered* *such that* $u_1 > u_2 > \ldots > u_{N-1} > l_{N-1} \geq l_{N-2} \geq \ldots l_1$*) such that*

$$\psi_1 = \begin{cases} h(s) & \textit{if } s < l_1 \textit{ or } s > u_1 \\ 0 & \textit{otherwise} \end{cases},$$

$$\psi_n = \begin{cases} h(s) & \textit{if } l_{n-1} < s < l_n \textit{ or } u_{n-1} > s > u_n \\ 0 & \textit{otherwise} \end{cases} \quad \textit{for } 1 < n < N,$$

$$\textit{and } \psi_N = \begin{cases} h(s) & \textit{if } s > l_{N-1} \textit{ or } s < u_{N-1} \\ 0 & \textit{otherwise} \end{cases}.$$

Figure 5 is an example of a potential optimal strategy. District 1 comprises a slice of extreme Republicans and a slice of extreme Democrats, and this slicing proceeds toward the center of the signal distribution. The slices from the right tail of the signal distribution contain more mass than the matched slice from the left tail, lest Republicans "cut it too close" in accounting for the noisy measurement of preferences. This follows the intuition developed in the third example in Section 3.

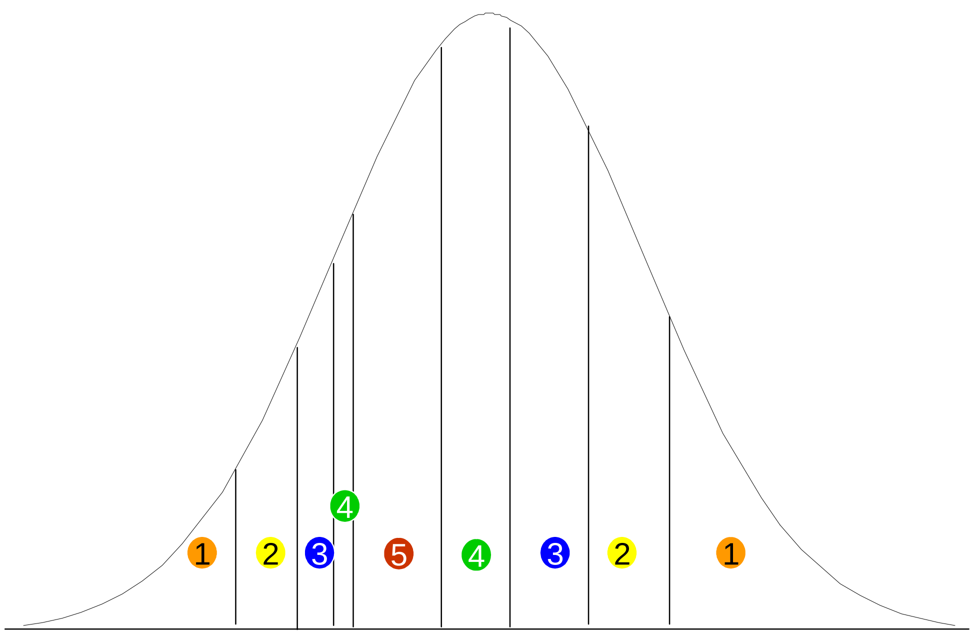

Figure 5: An example of the optimal strategy

4.4 Comparison with Received Literature

Previous work has considered two types of models which are both special cases of our model. The approach most similar to ours is that of Gilligan and Matsusaka (1999), in which voters always vote for a given party and their preferences are known with certainty to the gerrymanderer. Our model simplifies to this case (as shown in the first example in Section 3) if the conditional preference distribution limits to a point-mass at the true preference (so that preferences are observable) and if the breakpoint distribution $B(\cdot)$ is a point mass (so that voters are either Democrats or Republicans). As such, our model is more general and captures an important intuition - that more noise leads the gerrymanderer to create a larger buffer. Furthermore, our model has a continuum of preferences, and therefore is instructive not only as to the optimal number of Republicans and Democrats in a district but also which types of Republicans and Democrats should be combined.

The second approach to modelling gerrymandering - one perhaps more popular than that of Gilligan and Matsusaka - is a binary signal model with noise. In such a model (e.g. Owen and Grofman (1988)), the optimal strategy involves "throwing-away" some districts and "smoothing" others. Owen and Grofman refer to this as a "bipartisan gerrymander," since there are Democratic districts (those "thrown away") and Republican districts (the others). For instance if 60% of the population have signal r and 40% signal d, then the optimal strategy involves creating a certain number of districts which contain only those with signal d, and spreading the r voters uniformly over the remaining districts. This result is also a special case of our model, with additional assumptions, as shown in Proposition 2.

Proposition 2 *Suppose $s \in \{d, r\}$ and that Conditions 1 and 2 hold. Suppose further that $B(\cdot)$ is unimodal, with mode greater than d and less than r. Then the optimal gerrymander involves creating some districts with all voters of type d, and others with a constant proportion of r and d, and possibly one "odd district" with a non-zero but less-than-half proportion of r (from integer rounding problems). When $N \rightarrow \infty$, the optimal solution is a pure "bipartisan gerrymander."*

Thus our model nests the solution of "bipartisan gerrymandering," but the conclusions of such a model are very sensitive to several extreme assumptions. Furthermore, the intuitions which this special case highlights are very misleading. For instance, suppose that there are three signals: r, d and i (independents). As Proposition 1 shows, the optimal strategy matches increasingly extreme segments from the right and left tails (in this case Republicans and Democrats) into the same districts. The district which Republicans have the lowest chance of winning is not that which contains many Democrats, but rather one which contains many Independents. That is, these least Republican districts contain voters from the middle of the signal distribution, not the extreme left tail. It is also clear that "smoothing" is not a robust intuition. It is true only in the special case of a binary signal because there is no heterogeneity among potential Republican voters.

4.5 Comparative Statics

In this subsection we consider how the value to being the gerrymanderer responds to changes in the underlying distribution of voter preferences and signals. We also consider how this value changes as the number of districts to be created changes.

Our first comparative static shows that more precise signals are always better for the gerrymanderer.

Definition 1 *Consider two conditional preference distributions g and g'. g provides a More Precise signal than g' if there exists a conditional distribution $c(s' \mid s)$ such that*

$$\int g(\beta \mid s')c(s' \mid s)ds' = g\left(\beta \mid s\right).$$

Proposition 3 *The expected number of districts won by the gerrymanderer is increasing in the precision of the signal.*

This result shows that the gerrymanderer wins more districts in expectation as the signal received becomes more precise. Intuitively, as the gerrymanderer receives a better signal, the need for a large "buffer" of voters in a district declines. Instead, she can construct districts of a given median with a smaller proportion of voters from the right hand tail, leaving more right-wingers for other districts. Mathematically, the gerrymanderer could always lower the quality of the signal, while the reverse operation is not possible. Thus, it cannot be that a lower quality signal is better.

Our second comparative static result shows that the gerrymanderer does better as the distribution of voters becomes more spread out.

Proposition 4 *Consider two joint distributions $F(\beta, s)$ and $\hat{F}(\beta, s)$, with marginal distributions of β given by $F(\beta)$ and $\hat{F}(\beta)$, such that $\hat{F}(\beta)$ is a symmetric spread of $F(\beta)$. Then the expected number of districts won by the gerrymanderer is higher for $\hat{F}$ than for F.*

Intuitively, our model assumes that all signals have the same variance of preferences. But, if the breakpoint is more likely to be near the center of the preference distribution, then there is less uncertainty as to the voting patterns of extreme voters. For instance, suppose the breakpoint is normally distributed. If a voter has either $\beta = -0.5$ or $\beta = 0.5$, she will vote Republican either 31% or 69% of the time, quite a bit of uncertainty; but if a voter has either $\beta = 1.5$ or $\beta = 2.5$, she will vote Republican either 93% or 99% of the time. Extreme voters are thus more valuable to the gerrymanderer. Since an increase in the variance of the voter preference distribution increases the share of extreme voters in the population, the expected number of seats won increases.

The final comparative static concerns the number of districts.

Proposition 5 *Suppose that the number of districts increases by an integer multiple (that is doubles or triples). Then expected percentage of districts won by the gerrymanderer strictly increases.*

In previous analyses in this literature, proportional increases in the number of districts has little import; if twice the number of districts are required, the existing districts are split into equal parts, and so the voter profiles of the districts do not change. Our model implies that such parfaits are inefficient. Instead, the gerrymanderer can do better by slicing within previous districts, grouping the most and least Republican voters from an old district into one new district and giving the all less extreme voters to the other.

5 Numerical Examples

In order to illustrate the characterization of the optimal gerrymandering strategy and its comparative statics, we report the results of a number of numerical examples in this section. The examples all assume that there are five districts and that the gerrymanderer is Republican. In these examples we assume that the joint distribution of preferences and signals, $F(\beta, s)$ is multivariate normal with parameters

$\mu_\beta = \mu_s = 0$ and covariance matrix Σ with:

$$\Sigma = \begin{pmatrix} \sigma_\beta^2 & \rho\sigma_\beta\sigma_s \\ \rho\sigma_\beta\sigma_s & \sigma_s^2 \end{pmatrix}$$

. This implies that both the signal distribution and the conditional preference distribution are themselves normal. Note that this assumption satisfies Conditions 1 and 2. In this base case we assume a distribution of $F(\beta, s)$ such that $\beta \sim N(0, 5)$, $\rho = 0.5$. Furthermore, we assume that $\sigma_s = \rho\sigma_\beta$ so that $G(\beta \mid s) \sim N\left(s, \sigma_{\beta|s}^2 = \sigma_\beta^2 (1-\rho)\right)$. In all examples, we let $B \sim N(0,1)$ and set $N = 5$. Note that these assumptions imply that, nominally, half the voters are Republicans and half are Democrats - without gerrymandering, each party would win 2.5 seats, in expectation.

Table 1: Baseline Example

	District				
	1	2	3	4	5
Upper Slice	0.62	0.73	0.91	1	N/A
Lower Slice	0.38	0.27	0.09	0	N/A
Prob (win)	87.5%	74.8%	65.7%	41.7%	13.7%

Table 1 highlights a number of features of the optimal strategy. First, the highest median district (district 1) consists of 62% from a slice from the right tail of the distribution and 38% from a slice from the left tail. These upper slices get progressively larger for the lower median districts. While district four comprises a whole slice, districts 1 through 3 are formed by matching slices from the right and left tails. (District 5 consists of a whole slice containing those voters remaining after removing the first four districts from the signal distribution, and so the fraction in the upper and lower slice is not relevant). Second, note that the probability of winning district 1 is very high - 87.5%. This means that those in the left-most part of the distribution have very little chance of gaining representation. Third, no districts are "thrown away"; the gerrymanderer has more than a 13% chance of winning even the district least favorable to her. If she had "thrown-away" the district - that is, put those with the lowest signal into it - then, in this example, she would only win it 1.4% of the time. Finally, the number of districts won in expectation in this case is 2.8, compared with a non-gerrymandered equal representation benchmark of 2.5.

Hence, in this case the ability to be the gerrymanderer leads to a 13% increase in the expected number of districts won.

Table 2 illustrates how a change in the spread of the conditional preference distribution affects the gerrymanderer.

Table 2: Signal Coarseness

		District / Probability Won				
Signal	**Expected Districts Won**	1	2	3	4	5
$\sigma^2_{\beta\|s} = 0.5$	3.4612	97.4%	86.9%	74.3%	56.6%	30.9%
$\sigma^2_{\beta\|s} = 2.5$	2.8343	87.5%	74.8%	65.7%	41.7%	13.7%
$\sigma^2_{\beta\|s} = 4.5$	2.5349	68.2%	61.9%	55.7%	41.8%	25.9%

In accordance with our comparative static results, the gerrymanderer does worse as the quality of her signal deteriorates. This is reflected in a lower probability of winning each district, and hence a lower overall value to being the gerrymanderer. For instance, note that when the signal is very coarse, $\sigma^2_{\beta|s} = 4.5$, the gerrymanderer wins only 2.54 districts in expectation - barely more than the 2.50 won under proportional representation. Also, in the $\sigma^2_{\beta|s} = 0.5$ case the gerrymanderer has a 31% chance of winning district 5 - if she "threw it away" that would be just 0.2%. Finally, although the expected districts won, and hence the value function, is monotonic in $\sigma^2_{\beta|s}$ (as we have shown analytically), the probability of winning each district is not monotonic. Intuitively, as the signal becomes more informative the gerrymanderer can cut the districts finer, but the probability of winning the votes of those with the lowest signals decreases. These two effects work in opposite directions, which leads to the potential non-monotonicity of the probability of winning districts with "low" medians (here districts 4 and 5).

Table 3 shows how a change in the spread of the voter preferences affects the gerrymanderer.

Table 3: Spread of Voter Preferences

		District / Probability Won				
Signal	**Expected Districts Won**	1	2	3	4	5
$\sigma_\beta^2 = 3$	2.5528	71.0%	62.3%	55.6%	41.2%	25.1%
$\sigma_\beta^2 = 5$	2.8343	87.5%	74.8%	65.7%	41.7%	13.7%
$\sigma_\beta^2 = 25$	3.7802	100%	97.1%	90.6%	73.9%	16.4%

As voter preferences become more spread out, the gerrymanderer does better as our comparative static results showed. There is a monotonic increase in the probability of winning districts 1-4 as voter preferences become more spread out, since fewer extreme voters are necessary to provide a solid margin of victory (in expectation). A similar non-monotonicity as discussed above is at work here with the probability of winning district 5.

Table 4 reports how changes in the mean affect gerrymandering. A natural interpretation of a change in the mean is that it is a change in the number of nominal Republicans/Democrats. With the mean at zero there are 50% nominal Republicans. As the mean increases the share of nominal Republicans increases and *vice versa.*

Table 4: Nominal Republicans

			District / Probability Won				
% Repub.	**E[Win]**	**"Value"**	1	2	3	4	5
30%	2.0433	0.5790	49.4%	47.0%	40.7%	27.8%	10.2%
40%	2.4362	0.4837	87.0%	73.0%	52.3%	25.1%	6.2%
50%	2.8343	0.3343	87.5%	74.8%	65.7%	41.7%	13.7%
60%	3.2427	0.1951	87.8%	76.1%	67.3%	58.6%	34.5%
70%	3.6656	0.1199	90.2%	79.6%	71.7%	65.0%	59.1%

Note that as the proportion of nominal Republicans increases, both the expected number of seats won and the value to being the gerrymanderer increase. This value represents the difference in expected seats won compared to proportional representation.

6 Extensions

In this section, we briefly discuss the implications of some extensions to the basic model.

6.1 Majority Power and Risk Aversion

Our analysis considers a gerrymanderer whose payoff function is equal to the expected number of districts won. This is likely a good approximation for Congressional districting, where the uncertainty over the eventual party balance in the House of Representatives makes each district in a given state equally important.[13] But in state legislature districts, other objective functions may be important. For instance, a party might derive great benefit from remaining in the majority, in which case the gerrymanderer's value function would include a positive discontinuity at winning a majority of the districts. The marginal benefit to the gerrymanderer from each seat won might be diminishing as she wins more seats, it which case the objective function would become concave. The next proposition shows that, for any objective function whose argument is the percent of seats won in an election, our characterization still holds.

Proposition 6 *Suppose that the gerrymanderer designs districts to maximize*

$$E\left[V\left(\frac{1}{N}\sum_{n=1}^{N}d_n\right)\right],$$

where $d_n = 1$ *if the Republicans win district* n *and* 0 *otherwise and* V *is any strictly increasing function. Then Proposition 1 still characterizes the optimal partisan gerrymander.*

The impact of a change in the value function on the optimal redistricting plan will, in most cases, change the size of the upper and lower slices for many districts. It will not, however, change the form of the optimal partisan gerrymander, in which districts comprise increasing extreme slices from opposite tails of the signal distribution. Since district composition only affects the probabilities of various outcomes,

[13] One notable exception to this equality is the election of the president. Should the electoral college fail to produce a majority winner (as was the case in 1800, 1824, and 1876), the election would be thrown into the House of Representatives, where each state delegation receives a single vote.

the marginal benefit of increasing the probability of victory in any district (holding the composition of the other districts fixed) is a constant. Thus, Proposition 1 holds in this more general case just as in the simpler case discussed above where the objective function is simply the expected number of seats won.

Other changes in the optimal districting scheme depend on the balance of voters in the state and the shape of the value function. For instance, suppose the objective function were linear but for a positive discontinuity at winning a majority of the seats. Under normal circumstances, where the gerrymanderer possesses a commanding popular majority, state redistricters would be effectively risk-averse. As compared to the base case, the optimal plan would win fewer districts with greater certainty by including more Republicans in the right tail slice of the district to ensure remaining in the majority. But if the gerrymander faces a hostile population (perhaps due to the iniquities of gerrymanders past), the party would become risk-loving.

Another deviation from the value function analyzed above is for the gerrymanderer to be risk-averse, in the sense that they have a have a concave von Neumann-Morgenstern expected utility function. The clear implication of this is that the optimal strategy would be altered so that there was a larger "buffer" in districts which the gerrymanderer expects to win. A Republican gerrymanderer would put relatively more voters from the right tail of the signal distribution into the high median districts. This would raise the median of such districts, and lower the median of the low median districts. This altered strategy lowers both the expected number of districts won the variance of that total.

6.2 Voter Turnout

In our model we have implicitly assumed that everyone votes; obviously, in a system with non-compulsory voting, voter turnout is a real and important issue. Our results are unchanged if, for each type of voter, the turnout rate (which may vary across voter types) is independent of district composition. Furthermore, if turnout is statistically independent of the signal though not constant within voter classes, our results are unchanged. Our results are strengthened if moderate voters of a given nominal party affiliation are more likely to vote when matched with more extreme voters of the same party. In this case it is even more important to match extreme

Republicans with extreme Democrats, lest moderate Democrats turnout in great numbers. Endogenous voter turnout based on the nuances of district composition requires a model with additional structure.

6.3 Incumbent Effects

Another empirical regularity of Congressional races is the seemingly large electoral advantage enjoyed by incumbents - fewer than 3% of incumbents are defeated in the typical election cycle. There are three possible causes for this edge. First, an incumbent may simply reflect the preferences of her constituents. In this case, incumbency is simply a proxy for match quality between a Representative and her district, and one can say that incumbency, *per se*, has no effect on the conclusions of this model. Second, the incumbent may be more well known to her constituents in a variety of ways, and thus more easily elected; a (Republican) gerrymanderer would respond to this type of incumbent advantage by maintaining Republican incumbent districts as constant as possible, while matching Democratic incumbents to new and unfamiliar (though not necessarily different, from a signal profile perspective) districts. Indeed, such tactics were a key part of the Republican gerrymander of Texas in 2003. This effect is primarily a geographic concern, though, and is thus somewhat orthogonal to the predictions of our model.

A third source of advantage for an incumbent may be, broadly speaking, her resumé of Congressional experience and the resulting low quality of opponents, an edge which would follow her no matter the geographic specifics of her district. Ansolabehere, Snyder, and Stewart (2000) use the decennial redrawing of district boundaries to estimate that this third channel accounts for one-third to one-half of the incumbency advantage, on average, though there is surely much individual heterogeneity in the magnitude of the effect. The conclusions of our model would change in the presence of large incumbent effects of this third type, which would, in effect, make the distribution of the electoral breakpoint B district-specific. For instance, suppose that a particular Democratic incumbent was universally well liked and assured of election regardless of the composition of her district. It would then be optimal for a Republican gerrymanderer to "throw away" her district by including in it the most extreme democrats.

6.4 Endogenous Candidate Positioning

Endogenous location of politicians may also cause the distribution of electoral breakpoints to differ across districts. Intuitively, one would expect the B distribution to shift to the right in districts with more right-wing voters, as the Democratic candidates in such districts would be more conservative. The shifting breakpoint distribution would dampen the effect on the value function of increasing the size of the right-wing slice in a given district - but, so long as this effect were not larger than the direct effect of such a change, the optimal strategy would have the same characteristics.

Once candidates differ, the value function for the gerrymander might also change. This would serve to offset the effect of endogenous location because, though a Democrat might have a greater chance of election, she would be a more conservative democrat. In fact, in the limit of both cases where the gerrymander cares only about the positioning of the elected candidate and where both elected candidates converge to the district median voter (*à la* Hotelling (1929)), the value function would remain the same, as would the optimal strategy. Such a model has the further implication that, as gerrymandering becomes more pronounced (perhaps due to technological advances in recent years), Representatives to the House would become more extreme relative to Senators, who represent the median voters of each state. We empirically investigate this hypothesis in other work.

7 Conclusion

In this paper, we show that existing models of partisan gerrymandering make simplifying assumptions which have drastic implications for the conclusions which they draw. We analyze a substantially more general model with a continuum of voter preferences and noisy signals of those preferences. The model nests major models in the literature as special cases. The optimal strategy in our model creates districts which match extreme Republicans and extreme Democrats, rather than "throwing away" districts and "smoothing" over others. This characterization of the optimal partisan gerrymander is robust to a number of extensions, including alternative partisan objective functions, stochastic voter turnout, and endogenous candidate positioning.

The primary import of our paper is to suggest that existing models of partisan gerrymandering, and the intuitions behind them, are rather misleading. These intuitions are not simply academic speculations but rather are important in the world and give rise to conventional wisdoms about partisan gerrymandering which are inaccurate. For instance, traditional models imply that groups who have very different preferences from the gerrymanderer do not fare so badly - that is, although gerrymandering makes them worse off than proportional representation, they are assured of a lower bound of representation due to the gerrymanderer's "throwing away" some districts. Our model has very different implications. Instead, because of the "matching slices" strategy, they are combined into districts with a larger group of voters who have extremely different preferences from them, and so they have *very little* representation as a result of gerrymandering. Thus, our model suggests that the negative consequences of partisan gerrymandering for minority representation in government may be far worse than currently thought.

A natural question which follows from this analysis is to ask who are the voters in the opposite tail of the distribution to the gerrymanderers. We consider the empirical linkage between partisan and racial gerrymandering in related work. To illustrate this connection, suppose that the gerrymanderer is a Republican and that African-Americans are highly represented in the far left tail of the signal distribution (i.e. they have characteristics which make them very likely to vote for Democrats). In this case, under the optimal gerrymander, African-Americans would be placed in districts such that they receive very little representation. Data from the 2000 U.S. census and the 2000 presidential election suggests that African-Americans do indeed constitute the far left tail, and so an implementation of the optimal strategy, as characterized in this paper, would be severely disadvantageous to that population. The unmistakable implication of these facts is that partisan gerrymandering and racial gerrymandering are basically synonymous *in effect.* Since the 1960s, however, the Court has adopted a test based on intent, rather than effect.

A further implication of our analysis is that gerrymandering can be very valuable, and indeed is more valuable today than ever before. Technological advances have allowed gerrymanderers to gain better information about voters - in our model, a less coarse signal distribution in the sense of Blackwell - and draw boundaries with a finer pen. One would therefore expect parties to use an increasingly large amount of resources in order to become the gerrymanderer. Since the practice itself is probably lowers social welfare, spending resources on it merely exacerbates the social loss associated with gerrymandering. This implies that the welfare loss from gerrymandering is linked to such technologies, and has grown over time. In combination with the increasingly prevalent districting strategy of "kidnapping" - a practice also made easier with technology - we have little doubt that the problems of partisan gerrymandering are far worse today than at any time since *Baker*.

There are two clear directions for future work. The first involves empirical investigations of gerrymandering in the light of the theory developed here. The structure provided by our characterization of the optimal gerrymandering strategy is important for such empirical work. Previous empirical work on gerrymandering (see, for instance, Gelman and King (1990, 1994)) assumes a non-microfounded structural model which may give inaccurate estimates of the degree of gerrymandering. The second set of open issues involves the regulation of gerrymandering. Enriching the model to capture spatial considerations would make it possible to analyze the impact of constraints such as compactness. Although there is a body of work which attempts to deal with spatial considerations, the underlying models of gerrymandering which they employ are, as we have discussed, insufficiently rich to capture the core intuitions of the optimal strategy.

Ultimately, the effect of gerrymandering is an empirical question. As our model highlights, the impact of it depends on the particulars of the signal and preference distribution. However, one thing which this paper demonstrates is that empirical investigations alone can be misleading. Without understanding the optimal strategy for a gerrymanderer, one cannot properly assess the impact of partisan gerrymandering.

References

Ansolabehere, Stephen, James M. Snyder, and Charles Stewart III, "Old Voters, New Voters, and the Personal Vote: Using Redistricting to Measure the Incumbency Advantage," *American Journal of Political Science*, 2000, *44*, 17–34.

_ , _ , and _ , "Candidate Positioning in U.S. House Elections," *American Journal of Political Science*, 2001, *45*, 36–59.

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

Gelman, Andrew and Gary King, "Estimating the Electoral Consequences of Legislative Redistricting," *Journal of the American Statistical Association*, 1990, *85*, 274–282.

_ and _ , "Enhancing Democracy Through Legislative Redistricting," *American Political Science Review*, 1994, *88*, 541–559.

Gilligan, Thomas W. and John G. Matsusaka, "Structural Constraints on Partisan Bias under the Efficient Gerrymanderer," *Public Choice*, 1999, *100*, 65–84.

Hotelling, Harold, "Stability in Competition," *Economic Journal*, 1929, *39*, 41–57.

Issacharoff, Samuel, Pamela S. Karlan, and Richard H. Pildes, *The Law of Democracy: Legal Structure of the Political Process*, New York, NY: Foundation Press, 2002.

Karlin, Samuel and Herman Rubin, "The Theory of Decision Procedures for Distributions with Monotone Likelihood Ratio," *Annals of Mathematical Statistics*, 1956, *27*, 272–299.

Lee, David S., Enrico Moretti, and Matthew J. Butler, "Do Voters Affect or Elect Policies? Evidence from the U.S. House," *Quarterly Journal of Economics*, 2004, *119*, 807–859.

Milgrom, Paul R., "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 1981, *12*, 380–391.

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

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

Shotts, Kenneth W., "Gerrymandering, Legislative Composition, and National Policy Outcomes," *American Journal of Political Science*, 2002, *46*, 398–414.

8 Appendix

Proof of Lemma 1. The proof is by contradiction. The maximization problem can be described by the Lagrangian

$$L=\sum_{n=1}^{N}B\left(\mu_n\right)-\sum_{n=1}^{N}\lambda_n\left[\int_{-\infty}^{\infty}G\left(\beta\mid s\right)\psi_n\left(s\right)ds-\frac{1}{N}\right], \tag{4}$$

in addition to the boundary constraints. Suppose there exists a type of voter s with non-zero mass in each of two districts n and m and that $\mu_n \neq \mu_m$. In this case, it must be that the marginal value of that type of voter is equalized across district, which would imply that

$$b\left(\mu_n\right)\frac{\partial\mu_n}{\partial\psi_n\left(s\right)}-b\left(\mu_m\right)\frac{\partial\mu_m}{\partial\psi_m\left(s\right)}=0. \tag{5}$$

Implicitly differentiating (2), which determines the medians, yields

$$\begin{aligned}0 &= \int_{-\infty}^{\infty}g\left(\mu_n\mid s\right)\psi_n\left(s\right)ds\partial\mu_n+G\left(\mu_n\mid s\right)\partial\psi_m\left(s\right)\\ \frac{\partial\mu_n}{\partial\psi_n\left(s\right)} &= -\frac{G\left(\mu_n\mid s\right)}{\int_{-\infty}^{\infty}g\left(\mu_n\mid s\right)\psi_n\left(s\right)ds}=-\frac{G\left(\mu_n\mid s\right)}{\gamma_n\left(\mu_n\right)}.\end{aligned} \tag{6}$$

Equation (5) now implies that

$$\frac{b\left(\mu_n\right)}{\gamma_n\left(\mu_n\right)}G\left(\mu_n\mid s\right)=\frac{b\left(\mu_m\right)}{\gamma_m\left(\mu_m\right)}G\left(\mu_m\mid s\right)\Longrightarrow\frac{G\left(\mu_n\mid s\right)}{G\left(\mu_m\mid s\right)}=\frac{b\left(\mu_m\right)\gamma_n\left(\mu_n\right)}{b\left(\mu_n\right)\gamma_m\left(\mu_m\right)}. \tag{7}$$

Suppose now that there is another type of voters s' with positive mass in districts n and m. Without loss of generality, suppose that $s<s'$. This equation must hold for each point. Taking ratios of these equations yields that

$$\frac{G\left(\mu_n\mid s\right)}{G\left(\mu_m\mid s\right)}=\frac{G\left(\mu_n\mid s'\right)}{G\left(\mu_m\mid s'\right)}\Longrightarrow\frac{G\left(\mu_n\mid s'\right)}{G\left(\mu_n\mid s\right)}=\frac{G\left(\mu_m\mid s'\right)}{G\left(\mu_m\mid s\right)}. \tag{8}$$

Condition 1 implies that

$$\frac{G\left(\beta'\mid s'\right)}{G\left(\beta'\mid s\right)}<\frac{G\left(\beta\mid s'\right)}{G\left(\beta\mid s\right)},\forall s'>s,\beta'>\beta,$$

though, and since $\mu_n \neq \mu_m$, the equality in 8 violates Condition 1 - a contradiction.

$\blacksquare$

Proof of Lemma 2. Again, the maximization problem can be described by the Lagrangian in (4). Consider districts n and m, and suppose that $\mu_n < \mu_m$. The benefit of removing people of type s from district m and adding them to district n is

$$\phi_{mn}(s) = \frac{\partial \mu_n}{\partial \psi_n(s)} - \frac{\partial \mu_m}{\partial \psi_m(s)} = \frac{b(\mu_m)}{\gamma_m(\mu_m)} G(\mu_m \mid s) - \frac{b(\mu_n)}{\gamma_n(\mu_n)} G(\mu_n \mid s).$$

While this need not be positive for all s in district n, it must be, $\forall s \in \psi_m$ and $s' \in \psi_n$, that $\phi_{mn}(s') > \phi_{mn}(s)$. Note that $\phi'_{mn}(s) > 0$ is equivalent to $\frac{z(\mu_m|s)}{z(\mu_n|s)} < \frac{b(\mu_n)\gamma_m(\mu_m)}{b(\mu_m)\gamma_n(\mu_n)}$, and since the left-hand side is monotonically increasing in s from Condition 1, $\phi_{mn}(s)$ is quasi-concave. If $s'_1, s'_2 \in \psi_n$, then, for any point $s \in [s'_1, s'_2]$, $\phi_{mn}(s) > \min[\phi_{mn}(s_1), \phi_{mn}(s_2)]$. Thus $s \notin \psi_m$. $\blacksquare$

Proof of Lemma 3. Suppose, by way of contradiction, that there exist districts m and n such that $\mu_m = \mu_n$. Without loss of generality, suppose that the entire state consists of only districts m and n. By Lemma 1, if $\mu_m = \mu_n$ then $\gamma_m(\mu_m) = \gamma_n(\mu_n)$. By Lemmas 1 and 2, those voters in districts m and n must make up one or two complete vertical slices of $f(s)$. By the continuity of $G(\beta \mid s)$ and the density of the two aforementioned slices, there must exist four voter types $s_1 < s_2 < \mu_m < s_3 < s_4$ such that $G(\mu_m \mid s_1) - G(\mu_m \mid s_2) = G(\mu_m \mid s_3) - G(\mu_m \mid s_4)$ and $\psi_n(s_1) > 0$, $\psi_n(s_4) > 0$, $\psi_m(s_2) > 0$, and $\psi_m(s_3) > 0$. Furthermore, one can find such a quartet of voter types such that $G(\mu_m \mid s_3) > 0.5$ and $G(\mu_m \mid s_2) < 0.5$. In words, one district contains some of the inner type of voters, while the other district contains some of the more extreme types of voters relative to the district medians. Now consider a perturbation in which an equal "number" of voters of types s_1 and s_4 are transferred to district m from district n, and similarly voters of type s_2 and s_3 are transferred from district m to n. By construction, both μ_m and μ_n remain unchanged, as does the value function; but $\gamma_n(\beta)$ and $\gamma_m(\beta)$ have changed. By definition,

$$\frac{\partial \gamma_n(\mu_n)}{\partial \psi(s)} = g(\beta \mid s),$$

and so the derivative of $\gamma_n(\mu_n)$ for perturbations of this type is

$$\begin{aligned}\partial\gamma_n(\mu_n) &= \frac{\partial\gamma_n(\mu_n)}{\partial\psi(s_2)} - \frac{\partial\gamma_n(\mu_n)}{\partial\psi(s_1)} + \frac{\partial\gamma_n(\mu_n)}{\partial\psi(s_3)} - \frac{\partial\gamma_n(\mu_n)}{\partial\psi(s_4)} \\ &= g(\mu_n \mid s_2) - g(\mu_n \mid s_1) + g(\mu_n \mid s_3) - g(\mu_n \mid s_4).\end{aligned}$$

But, by Condition 2, the modes of the lower signals lie below μ_n. Thus, we know that $g(\mu_n \mid s_2) > g(\mu_n \mid s_1)$, and similarly that $g(\mu_n \mid s_3) > g(\mu_n \mid s_4)$, and so $\partial\gamma_n(\mu_n) > 0$. By likewise reasoning, $\partial\gamma_m(\mu_m) > 0$. After performing such a perturbation the new districting arrangement now violates Lemma 1 since $\mu_m = \mu_n$ while $\gamma_n(\beta) \neq \gamma_m(\beta)$. This new arrangement is not optimal, but the value function is unchanged from the old districting plan, and so the old plan cannot be optimal either - a contradiction. ■

Proof of Lemma 4. Suppose, by way of contradiction, that such a case existed. Without loss of generality, from Lemma 2, we can assume that districts n and p comprise one whole slice. It also must be that $s^* < s'$ for all $s' \in \psi_n$ and than $s^* > s''$ for all $s'' \in \psi_p$. Denote $\bar{s}_n = \sup\{s \in \psi_n\}$, $\bar{s}_p = \sup\{s \in \psi_p\}$, $\underline{s}_n = \inf\{s \in \psi_n\}$, and $\underline{s}_p = \inf\{s \in \psi_p\}$. Of course, $\bar{s}_n > \underline{s}_n > s_m > \bar{s}_p > \underline{s}_p$.

The Lagrangian from equation 4 implies that, if $s \in \psi_j$, then

$$-a_j G(\mu_j \mid s) - \lambda_j \geq \max_i -a_i G(\mu_i \mid s) - \lambda_i$$

where $a_i = \frac{b(\mu_i)}{\gamma_i(\mu_i)}$. These a_i coefficients represent the sensitivity of the median of district i to changes. For each district i, denote these expressions by η_i. Since it must be, by Lemma 2, that $\lim_{s\to\infty}\eta_m(s) > \lim_{s\to\infty}\eta_n(s)$, we know that $\lambda_m < \lambda_n$. By a similar argument, it must be that $\lambda_n < \lambda_p$. We also know that

$$\eta_n(\bar{s}_n) \geq \eta_m(\bar{s}_n) \quad \text{and} \quad \eta_m(s^*) \geq \eta_n(s^*),$$

which implies that

$$a_m \leq a_n \frac{G(\mu_n \mid s^*) - G(\mu_n \mid \bar{s}_n)}{G(\mu_m \mid s^*) - G(\mu_m \mid \bar{s}_n)}. \tag{9}$$

(9) states that district m must not be too sensitive compared to district n. Were this so, a profitable deviation would exist by shifting district n down to include s^*

and giving voters of type $\bar{s}_n$ to district m. Similar arguments imply that

$$a_m \geq a_p \frac{G\left(\mu_p \mid \underline{s}_p\right) - G\left(\mu_p \mid s^*\right)}{G\left(\mu_m \mid \underline{s}_p\right) - G\left(\mu_m \mid s^*\right)}, \tag{10}$$

which has the interpretation that district m must be sensitive enough relative to district p so that shifting district p up to include s^* is not profitable. Note that the right-hand side of the inequality in (10) is greater than 1 by the unimodality of $G(\beta \mid s)$ and by the fact that $\mu_p \in \left(\underline{s}_p, s^*\right)$. Of course, (9) and (10) can only hold simultaneously if the right-hand side of (9) is greater than or equal to the right-hand side of (10). This requires

$$\frac{a_p}{a_n} = \frac{b\left(\mu_p\right)\gamma_n\left(\mu_n\right)}{b\left(\mu_n\right)\gamma_p\left(\mu_p\right)} \leq \frac{G\left(\mu_n \mid s^*\right) - G\left(\mu_n \mid \bar{s}_n\right)}{G\left(\mu_p \mid \underline{s}_p\right) - G\left(\mu_p \mid s^*\right)} \frac{G\left(\mu_m \mid \underline{s}_p\right) - G\left(\mu_m \mid s^*\right)}{G\left(\mu_m \mid s^*\right) - G\left(\mu_m \mid \bar{s}_n\right)}. \tag{11}$$

Now consider what happens to this ratio as we increase the precision of the signal (which can be thought of here as shrinking the conditional preference distribution G into the median). Since district n contains voters closer in signal to the median of district m, the ratio $\frac{G\left(\mu_m \mid \underline{s}_p\right) - G\left(\mu_m \mid s^*\right)}{G\left(\mu_m \mid s^*\right) - G\left(\mu_m \mid \bar{s}_n\right)}$ will shrink, going to 0 in the limit. On the other hand, both $G\left(\mu_n \mid s^*\right) - G\left(\mu_n \mid \bar{s}_n\right)$ and $G\left(\mu_p \mid \underline{s}_p\right) - G\left(\mu_p \mid s^*\right)$ rise to 1, since $\underline{s}_p < \mu_p < s^* < \mu_n < \bar{s}_n$. Thus, the right-hand side of (11) shrinks to 0 as the precision of the signal increases. Note, however, that the ratio $\frac{a_p}{a_n}$ is bounded away from 0, since $\frac{\gamma_n(\mu_n)}{\gamma_p(\mu_p)}$ will limit to 1 (by the definition of $\gamma(\mu)$) and $\frac{b(\mu_p)}{b(\mu_n)}$ is bounded away from 0 since the medians μ_n and μ_p are bounded and the c.d.f. B is strictly increasing. Thus, for sufficiently high signal quality, the inequality in (11) cannot hold - a contradiction. ■

Proof of Proposition 1. Apply Lemmas 1-4. ■

Proof of Proposition 2. Suppose not. The choice variable for each district can be summarized by ψ_n, the proportion of R in the district. Then there exist two districts m and n such that $\psi_m \neq \psi_n$ and $\psi_i > 0$ for $i = \{m, n\}$. Without loss of generality, let $\mu_m > \mu_n$. By Condition 1, $G(\beta \mid r)$ first order stochastically dominates $G(\beta \mid d)$, and so $\psi_m > \psi_n$.

In order that there be no profitable deviations, it must that $\frac{\partial \mu_n}{\partial \psi_n} = \frac{\partial \mu_m}{\partial \psi_m}$. But, in general,

$$\frac{\partial^2\mu}{\partial\psi_i^2}=\frac{\frac{\partial\mu_i}{\partial\psi_i}}{\gamma(\mu)^2}\left\{\begin{array}{c}[[g(\mu\mid d)-g(\mu\mid r)]\,b(\mu)+[G(\mu\mid d)-G(\mu\mid r)]\,b'(\mu)]\\ \cdot[\psi_i(g(\mu\mid r)-g(\mu\mid d))+g(\mu\mid d)]\\ -b(\mu)[G(\mu\mid d)-G(\mu\mid r)][\gamma'(\mu)+g(\mu\mid r)-g(\mu\mid d)]\end{array}\right\},$$

which is positive when $\mu<0$ and negative when $\mu>0$. Since $\mu>0\iff\psi>0.5$, the concavity of μ implies that one could never have $\psi_m>\psi_n\geq 0.5$, since then $\frac{\partial\mu_n}{\partial\psi_n}>\frac{\partial\mu_m}{\partial\psi_m}$, and so R could do better by increasing n and decreasing m. It also implies that there cannot be $0.5>\psi_m\geq\psi_n$, since then $\frac{\partial\mu_n}{\partial\psi_n}<\frac{\partial\mu_m}{\partial\psi_m}$ and the opposite deviation would improve R's representation. Thus, there can only be one "odd district" with $0<\psi<0.5$, and all districts with $\psi>0.5$ must have equal proportions of r and d.

Suppose that $N\to\infty$. Note that there can only be one "odd district." Let the mass of voters in this district have Lebesgue measure τ. Since each district must have an equal mass of voters, $\tau=1/N$. Clearly, $\lim_{N\to\infty}\tau=0$. ■

Proof of Proposition 3. First note that signal precision provides a partial ordering on conditional preference distribution. Now, if the signal contains no information then the expected number of seats won by the gerrymanderer is the population share. If the signal is perfectly precise such that $s=\beta$ it is possible (see Proposition 1) to create districts such that only the lowest median district has a median equal to the population median, while all other lie above. Hence the gerrymanderer wins more seats in expectation with a perfect signal. Now consider any two conditional preference distributions g and g' such that g provides a more precise signal than g'. The gerrymanderer must win at least as many seats in expectation under g than g' since the value function has the Blackwell Property. That is, she could construct a distribution c such that from g she could generate g'. ■

Proof of Proposition 4. Fix the optimal districting plan under $F(\beta,s)$ and consider the construction of the highest median district (wlog district 1) with median μ_1 given by $\int_{s\in\psi_1}G(\mu_1\mid s)h(s)ds=\frac{1}{2N}$, comprising an upper and lower slice. Let the upper slice contain w_1 share of the voters in the district. Suppose that, under $\hat{F}(\beta,s)$, the gerrymanderer sets $\hat{\mu}_1=\mu_1$. This can be achieved with at least as small an upper slice $\hat{w}_1\leq w_1$, since the Republican voters (who make up more than half of the district) are at least as likely to vote Republican as before. If $\hat{w}_1 < w_1$, then note that all other districts $2, ..., N$ have a higher medians even if we set $\hat{w}_n = w_n$ for all n, that is without re-optimizing their construction. If $\hat{w}_1 = w_1$, then repeat this procedure until finding a district n^* such that $\hat{w}_{n^*} < w_{n^*}$. By assumption that $\hat{F}$ has greater symmetric spread than F, this must be true for at least one district. Hence the value function under $\hat{F}(\beta, s)$ is higher than under $F(\beta, s)$. This reasoning must hold for any such pair of distributions. ■

Proof of Proposition 5. Consider an increase from N districts to mN, where m is an integer. By replication, the gerrymanderer could do at least as well with mN districts as with N - but this replication involves creating parfaits. From Lemma 3 this is a suboptimal strategy. Hence the value function under the optimal strategy must be higher. This completes the proof. ■

Proof of Proposition 6. Fix the districting scheme and consider the marginal benefit from a small deviation x, which would be

$$\frac{\partial E\left[V\right]}{\partial x} = b\left(\mu_n\right)\left(K_i - L_i\right)\frac{\partial \mu_n}{\partial x},$$

where K is the expected value if the Republican candidate wins in district n and L is that value if the Democrat wins in district n. Note that this is identical to the value derived from equation (5) but for the term $(K_i - L_i)$, which is fixed for all deviations from a districting plan. Thus, the "sensitivities" $\{a_i\}_{i=1}^{N}$ are now differently scaled, but the constant does not affect any proofs. Propositions 1 through 5 hold. ■

9 Figures

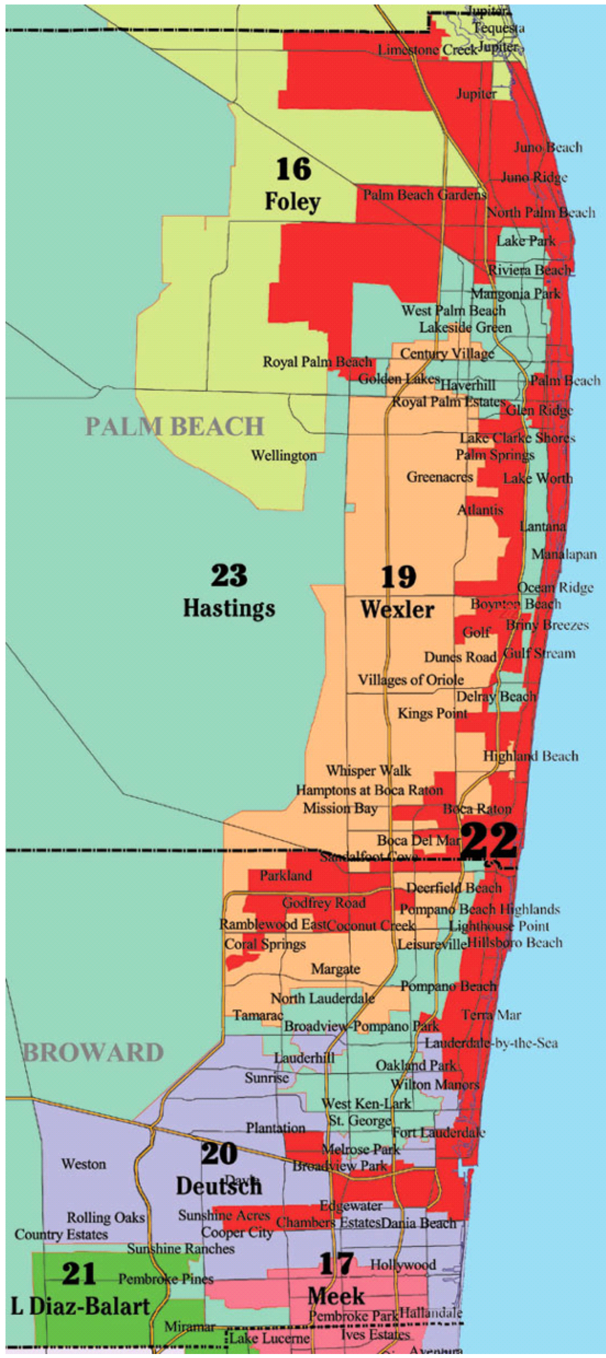

Figure 1: Florida 16th to 23rd Congressional Districts

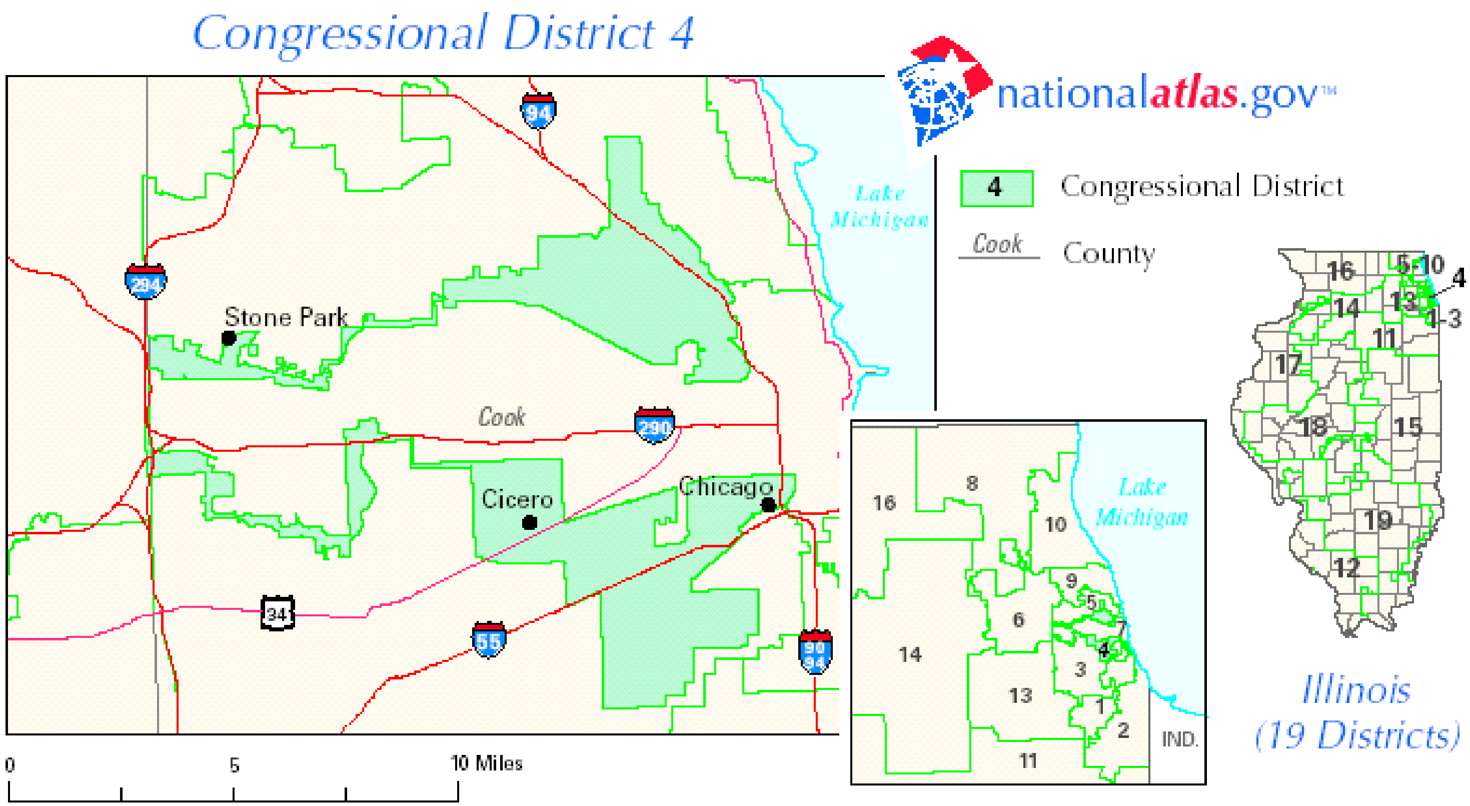

Figure 2: Illinois 4th Congressional District