A Statistical Test for Partisan Gerrymandering*

Preliminary Work. DO NOT CITE!

Scott T. Macdonell[†]

Abstract

Is redistricting the result of partisan gerrymandering or apolitical considerations? I develop a statistical test for partisan gerrymandering and apply it to the U.S. Congressional Districting plan chosen by the Republican legislature in Pennsylvania in 2001. First, I formally model the optimization problem faced by a strategic Republican redistricter and characterize the theoretically optimal solution. I then estimate the likelihood a district is represented by a Republican, conditional on district demographics. This estimate allows me to determine the value of the gerrymanderer's objective function under any districting plan. Next, I use a geographic representation of the state to randomly generate a large sample of legally valid plans. Finally, I calculate the estimated value of a strategic Republican redistricter's objective function under each of the sample plans and under the actual plan chosen by Republicans. When controlling for incumbency the formal test shows that the Republicans' plan was a partisan gerrymander.

1   Introduction

The United States adds an unusual wrinkle to the standard form of representative democracy; every ten years politicians are required to choose voters. Specifically, I consider the decennial redrawing of U.S. Congressional Districts by state

---

[†]The University of Texas at Austin; 1 University Station, C3100; Austin, TX 78712; E-mail: smacdonell@utexas.edu; web: http://www.scott-macdonell.com

Tab. 1: 5 District State with 440 Democrats and 560 Republicans

District	Democrats	Republicans	District	Democrats	Republicans
1	0	200	1	88	112
2	110	90	2	88	112
3	110	90	3	88	112
4	110	90	4	88	112
5	110	90	5	88	112

(a) "Democratic" Plan   (b) "Republican" Plan

legislatures. This process is designed to avoid an anti-majoritarian problem: Given varying growth rates across the U.S., if Congressional Districts were not occasionally redrawn, some Congresspeople could represent districts of a few thousand individuals, with others representing districts of a few million. In Reynolds v. Sims (1964) the U.S. Supreme Court decided that the constitution requires occasional redistricting to ensure that "one man's vote in a congressional election is to be worth as much as another's." However, the constitution makes no mention of who should draw new maps; this process has generally been left to state legislatures.

This fact gives incumbent politicians substantial power to allocate seats in Congress according to their own interests, possibly against the will of the majority. Their only universal legal constraints are that all districts must be contiguous and (approximately) equipopulous.[1] As a quick example of the power of gerrymandering, consider a five district state with 440 voters who always vote for Democrats and 560 voters who always vote for Republicans. Ignoring contiguity, two possible plans are shown in Table 1. The first plan has one district of 0 Democrats and 200 Republicans, and the other four districts have 110 Democrats and 90 Republicans. In this case Democrats would win an 80% super-majority of the seats while making up a minority of the population. On the other hand, the second plan may initially seem fair as it is "proportional;" Each district has 88 Democrats and 112 Republicans. However, in this case Democrats from the state would have no representation despite making up nearly half the population. As evidenced by this simple example, the choice of districts can significantly impact the allocation of seats in Congress.

After the 2000 Census, both houses of the Pennsylvania state legislature and

---

[1] Even the contiguity constraint must sometimes be relaxed if, for example, the state has some small islands. Also, different states may allow different margins of error on the equipopulous constraint.

the Governor's mansion were controlled by Republicans. This presented them with the opportunity to draw Pennsylvania's 19 Congressional Districts so as to increase the number of Republicans in the U.S. House of Representatives. In fact, in the 2002 Congressional elections, Republicans won 12 of the state's 19 seats, despite the fact that Pennsylvania seemed to be a swing state that leaned Democratic.[2] This lead several Pennsylvania Democrats to claim their rights to equal representation had been violated and mount a legal challenge to Pennsylvania's districting scheme. This culminated in their case, Vieth v. Jubelirer (2004), reaching the U.S. Supreme Court. The court declined to intervene, deciding that partisan gerrymandering cases were non-justiciable. In Justice Kennedy's controlling opinion he noted that there was no test to appropriately determine if a districting scheme was an unconstitutional attempt to deny members of one party representation, or one based on other, less invidious considerations.[3] However, he left open the possibility that such a test could be developed in the future and that at such a time it could be appropriate for courts to intervene in partisan gerrymandering cases.[4]

The purpose of this paper is to develop a statistical test to evaluate claims of partisan gerrymandering. I start by setting up the theoretical optimization problem implied by a claim that a particular plan is a partisan gerrymander: the party in control chose district demographics in order to maximize the expected number of Representatives from their party. Next, in order to calculate the value of the objective function under particular districting schemes, I estimate the relationship between district demographics and the probability of electing a Representative from a given party. Using precinct and census tract data from Pennsylvania in 2000 and GIS techniques, I then randomly generate a large sample of alternative districting schemes that respect contiguity and population equality; this method allows me to calculate the demographics of each district in each sample plan.

The final step of the test compares the estimated value of the objective

---

[2] Two years earlier, Democratic presidential candidate Al Gore won Pennsylvania by a larger margin than he won the national popular vote (which he won).

[3] Of course, the justice phrased it differently: "Because there are yet no agreed upon substantive principles of fairness in districting, we have no basis on which to define clear, manageable, and politically neutral standards for measuring the particular burden a given partisan classification imposes on representational rights. Suitable standards for measuring this burden, however, are critical to our intervention." (Justice Kennedy's Concurrence from Vieth v. Jubelirer (2004))

[4] This opinion was decisive as four justices wished to strike down the Pennsylvania districting plan as an unconstitutional partisan gerrymander while the other four wished to declare partisan gerrymandering cases non-justiciable essentially in perpetuity.

function under the actual districting scheme and under the sample schemes in order to test a "no partisan gerrymandering" null hypothesis. Formally, my null hypothesis is that partisan considerations were not used to develop the districting scheme. I will use plans generated by my algorithm to test this claim in the following manner: If the chances of a disinterested cartographer producing a scheme so favorable to the redistricting party were less than 5%, I then would reject the null hypothesis at the 5% level. This forces acceptance of the alternative, that partisan considerations were used to develop the plan for Pennsylvania. This would imply that charges of partisan gerrymandering are valid by definition.

There is a large theoretical and empirical literature on strategic redistricting, though it abstracts away from the geographic nature of the problem. Many papers set up and solve the optimization problem or gerrymandering game faced by the party in charge of redistricting (e.g. Friedman and Holden (2008); Gul and Pesendorfer (2010)) or attempt to find socially optimal rules which could be imposed on the redistricting process (e.g. Coate and Knight (2007)) Alternatively, many papers take redistricting schemes as given and estimate theoretical quantities like "bias"[5] in order to determine how much a particular plan favors one party (e.g. Cox and Katz (1999); King and Browning (1987)).

Ignoring the geographic nature of redistricting limits the applicability of these theoretical and empirical lines of research. As an example, consider a hypothetical state that leans slightly Democratic overall, but only because of one very Democratic urban area. Perhaps the rest of the state leans slightly Republican. An apolitical cartographer might create a plan with a few districts which include parts of the urban area, but with most districts only covering other parts of the state. Under this plan Democrats would generally win with large majorities in the few urban districts, while Republicans may generally win the rest of the districts more narrowly. This plan might look "biased;" in this slightly Democratic state the Republicans could expect to win more seats. However, such plans might be quite likely even using non-partisan districting principles. One would expect far different outcomes in a more homogeneous state. Clearly, the geographic distribution of voters has important implications for the study of redistricting. To a partisan gerrymanderer, the geographic distribution of voter demographics constrains the set of demographics possible

---

[5] Bias is often defined based on the share of the seats a party could expect to win if they earned half of the statewide vote. If the party should expect to win less than half the seats in this situation, then the plan is biased against them.

for Congressional Districts.

A state could be thought of as a collection of a large number of "small" areas (e.g. street addresses, census blocks or tracts, etc...). Redistricting can then be thought of as the process of aggregating those smaller units into a certain number of larger units (Congressional Districts), with the requirement that the larger units be equipopulous and contiguous. Figure 1.1 demonstrates this approach. Figure 1.1a breaks the state of Pennsylvania down into individual census tracts, while Figure 1.1b shows the actual 2002 Congressional District boundaries.

An alternative line of research heavily focuses on this geographic interpretation of the problem, but often ignores the important theoretical and empirical results. For instance, papers such as Garfinkel and Nemhauser (1970) and Rossiter and Johnston (1981) attempt to identify all possible contiguous and equipopulous solutions to particular redistricting problems, in the hopes of choosing the "best." However, this approach is only computationally feasible when the "small" units are in fact quite large. For example Garfinkel and Nemhauser's (1970) method uses counties as their small units and fails for a state with as few as 55 counties. This negates the possibility of using the possibly significant variation of within county demographic differences to increase the demographic variation between Congressional Districts.[6] Similar to this paper, Engstrom and Wildgen (1977) and Cirincione, Darling and O'Rourke (2000) take the alternative approach of randomly generating a large number of contiguous and equipopulous redistricting plans in order to evaluate claims that a particular state unconstitutionally used race as a predominant factor during redistricting.

This paper provides the first empirical test of theoretical partisan gerrymandering predictions using a single redistricting plan.[7] Using theoretical and empirical results, I design the test based on a random sample of geographically allowable plans. While scholars have previously used random redistricting methods to test racial gerrymandering claims, they've arbitrarily chosen their test statistics. For instance, Cirincione, Darling and O'Rourke (2000) simply count the number of majority-minority districts in each plan (actual and randomly

---

[6] Increases in computational power provide little hope of redeeming this approach. Altman and McDonald (2011) suggest that even for the modestly sized Wisconsin, when using census blocks, the number of potential districting schemes these methods would have to check may be on the order of the number of quarks in the universe.

[7] Others such as Gelman and King (1994) empirically test partisan gerrymandering predictions. However, they require a large number of redistricting plans in their analysis. Courts need the capability to evaluate a claim that an individual plan is a partisan gerrymander.

Fig. 1.1: Pennsylvania

(a) Census Tracts
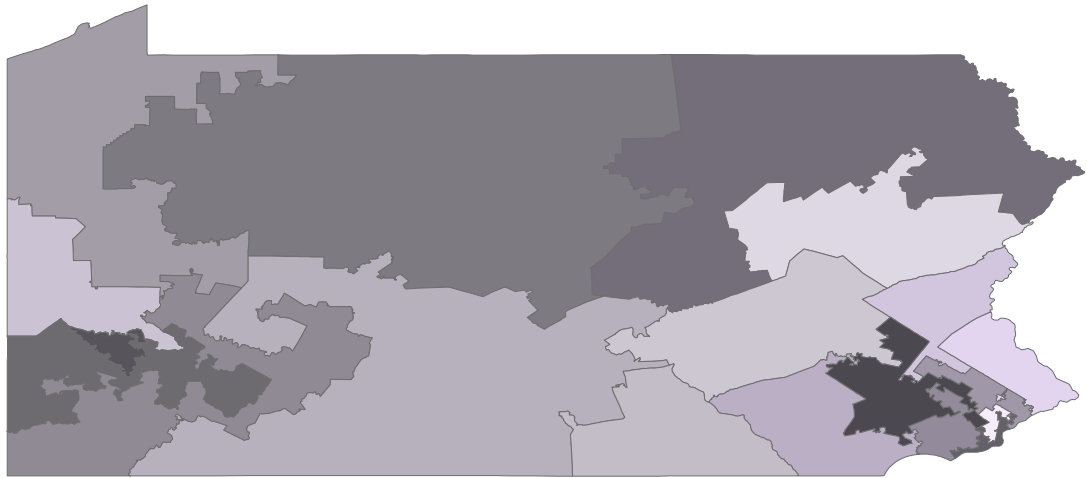
(b) 2002 Congressional Districts

generated).[8] This ignores the possibility that there is a range of population levels where minority representation is likely, yet uncertain.[9] Given these arbitrary test statistics, it may not be surprising that tests based on random samples of possible plans have not yet been used to evaluate claims of partisan gerrymandering. Do we simply count the number of "Democratic" districts? Is a district Democratic if it contains more registered Democrats? What if there are many independent voters?... I address this issue by formally modeling the optimization problem faced by a partisan redistricter and using the estimated value of their objective function as my test statistic.

My approach has the added value of being agnostic towards questions of fairness. For instance, Coate and Knight (2007) and Dopp (2011) try and determine redistricting principles that would induce more "optimal" districting schemes. Justice Kennedy was quite clear in his controlling opinion in Vieth v. Jubelirer (2004) that choosing among such principles was essentially a political question; any principle would likely favor one party over the other. Here, my test is not based on fairness or optimality, but on how abnormally beneficial a plan is to a particular party. Courts could avoid ruling based on whether a plan was fair or unfair. Instead, if the chances of a non-partisan cartographer producing a plan so favorable to the redistricting party were remote, courts could judge the plan a partisan gerrymander.

The remainder of this paper proceeds as follows: Section 2 formally defines the partisan gerrymanderer's optimization problem and characterizes the theoretically optimal solution ignoring geography. Section 3.1 discusses some of the data and estimates the parameters of the objective function. Section 3.2 describes the geographic data I use and the algorithm which generates my random sample of alternative plans. Section 3.3 provides the results. Section 4 further discusses the results of my analysis, explores some of the limitations of my approach, and examines further directions for related work. Section 5 concludes.

---

[8] However, they do explore the implications of a few alternative thresholds for when a district should be considered majority-minority.

[9] Engstrom and Wildgen (1977) begin to address this concern. Their test statistic assigns a 1 to each district in a plan where less than 45% of the population is a minority, a 3 in each district where minorities make up more than 55% of the population and the appropriate weighted average for districts in between. However, they do not provide a model or empirical estimates to justify this particular choice.

2   Theory

Partisan gerrymandering of U.S. House districts is generally assumed to be done with the aim of increasing the number of Representatives from the redistricter's own party in Congress. A partisan gerrymanderer's task is then to divide a state into N Congressional Districts in such a manner as to maximize the expected number of districts which elect a representative of her party. She may achieve her objectives by choosing the demographics for each district.

Without loss of generality, assume that the state is being gerrymandered by Republicans. Let x_i be a vector denoting the demographics of interest in district i. Suppose an element of this vector is a scalar "Republicaness" measure, r_i. I will discuss the interpretation and definition of r_i further in Section 3.1.

For now, assume that the only constraints the redistricters face are that each resident of the state must be in exactly one district, and that the districts must be of equal population. Let x (r) be the value of the demographics ("Republicaness") for the whole state. It can be shown that this implies the alternative constraint $\frac{\sum_{i=1}^{N} x_i}{N} = x$.[10] Of course, geography and correlation across demographics types add additional constraints, but ignore those until Section 3.2.

Let y_i represent the outcome of the upcoming election in district i. Define $y_i \equiv 1$ if district i elects a Republican, and $y_i \equiv 0$ if it elects a Democrat. Assume that there exists some known function, $F(x_i) \equiv Prob(y_{it} = 1 | x_i)$, which gives the probability of electing a Republican given district demographics. Then the gerrymanderer's optimization problem is as follows:

$$\max_{\{x_i\}_{i=1}^{N}} \sum_{i=1}^{N} F(x_i) \ s.t. \tag{2.1}$$

$$\frac{\sum_{i=1}^{N} x_i}{N} = x, \quad \underline{x} \leq x_i \leq \overline{x}, \ \forall i \tag{2.2}$$

where $\underline{x}$ and $\overline{x}$ represent the lower and upper bounds of the demographics. These could simply reflect the fact that it is impossible to make districts that are more than 100% (less than 0%) of a certain demographic type, or one could be

---

[10] So long as we suppose x measures the percentage of voters who are of various demographic types. For example, suppose one of those types is Hispanics. Given a particular plan, if we want to increase the portion of district one that is Hispanic by 1% we must remove 1% of the population from the district (choosing all non-Hispanics) and exchange them with Hispanics from other districts (by the population equality constraint). This process would keep the average percent Hispanic constant across the districts.

more restrictive with the bounds in an attempt to capture some yet unmodeled constraints faced by redistricters, such as contiguity.[11] Assume the problem is non-trivial for all demographics of interest $(\underline{x}, \ll x \ll \overline{x})$.

For the purposes of my statistical test, I could end this section with only equation 2.1. All I need is an objective function to evaluate under each districting scheme I consider. However, it is instructive to consider what a solution to this "geography-blind" optimization problem should look like. Comparing the solution to this version of the gerrymandering problem with the chosen plan provides a useful demonstration of the importance of the yet unmodeled geographic constraints.

Assume that r_i is the only demographic of interest which can vary across districts (x_i is a scalar equal to r_i).[12] Assume that $F(\cdot)$ is a smooth, symmetric, S shaped distribution.[13] Then, the following theorem applies and will allow me to find the optimal solution given a specific parameterization of the optimization problem:

Theorem 1. *For some $m \in \{0, 1, 2, ..., N-1\}$ one of the following two demographic profiles will be a solution to equations 2.1 and 2.2:*

1. $r_i = \underline{r}$ $\forall i \in \{1, ..., m\}$ (if $m = 0$, then this holds for none of the districts) and $r_j = q$ $\forall j \in \{m+1, m+2, ..., N\}$ where q solves $\frac{(N-m)q + m\underline{r}}{N} = r$.

2. $r_i = \underline{r}$ $\forall i \in \{1, ..., m\}$ and $r_j = \overline{r}$ $\forall j \in \{m+2, m+3, ..., N\}$ and $r_{m+1} = p$ where p solves $\frac{m\underline{r} + (N-m-1)\overline{r} + p}{N} = r$.

For a proof of this Theorem please see Appendix A. Once the optimization problem has been parameterized, this theorem takes an N dimensional optimization problem and allows one to find a solution by checking no more than $2N$ candidate solutions.[14] Furthermore, this Theorem mirrors the "pack and crack" result already common in the literature.[15] The Republican redistricters will "pack" many of their opponents into districts where the Democrats will generally win with overwhelming majorities. However, in the the rest of the districts the redistricter will spread the population more evenly (crack) such that

---

[11] It is more appropriate to simply include the geographic constraints as in Section 3.2.

[12] I will justify this assumption in Section 3.1.2.

[13] Formally, assume $\exists z \in R$ s.t. $\forall y > 0$, $1 - F(z+y) = F(z-y)$ and that $\forall w < z$, $F''(w) > 0$ and that $F(\cdot)$ is continuously differentiable.

[14] Some of the demographic profiles that satisfy Theorem 1 may not satisfy the constraints in equation 2.2. Therefore, one also needs to check the feasibility of each candidate solution under this Theorem.

[15] For a discussion of this common result as well as an example of the rare paper which finds conflicting results see Friedman and Holden (2008)

the Republicans have more moderate majorities. The logic is that the Republicans are certain to lose a few districts, but with many Democrats excluded from the remaining districts the Republicans are likely to capture a large majority of those.

3 Statistical Test

3.1 Estimation of $F(\cdot)$

In order to perform my statistical test I need to be able to evaluate equation 2.1 for any potential districting scheme. This will require that I be able to determine $\{x_i\}_{i=1}^{N}$ under any potential districting scheme, and that I estimate $F(\cdot)$. First, I discuss the dataset I use in this subsection. Then, I formalize my statistical model and discuss results.

3.1.1 Data

In order to complete the estimation stage of my analysis I used data from a variety of sources. First, I have presidential votes by Congressional District for 1972-2008 provided generously by Sean Theriault. This allows me to calculate a useful measure of "Republicaness." Specifically, I assume that r_i for any Congressional District is equal to the Republican presidential candidate's share of the major party vote in the most recent presidential election in district i minus the same share nationwide.[16]

Using this measure confers a number of advantages. For one, it directly measures how much more or less "Republican" a particular area is relative to the rest of the country. This leads to it being consistent across time. Estimating $F(\cdot)$ using multiple elections cycles and other demographics one might have to worry about a particular group's allegiance to each party changing over time. It seems less likely that there would be a significant change in how people who prefer Republican presidential candidates feel about Republican Congressional candidates. Furthermore, it inherently controls for the possibility that in the most recent election the Republican candidate may have been especially (un)appealing relative to his opponent.

---

[16] For example, suppose that nationwide in 2008 John McCain received 45% of the votes that were cast for either himself or Barack Obama. Also, suppose that in district i John McCain got 40% of the votes that were cast for either himself or Barack Obama. Then $r_i = -0.05$ for district i in 2008 and 2010.

Also, I have all U.S. House of Representatives election outcomes for 1972-1992 from ICPSR study 6311 (King, 2006).[17] This dataset also contains information on the presence of any incumbents, and their party.[18] Additionally, I have a large number of other demographic variables by Congressional District from 1972-1994 coming from a dataset produced by David Lublin.[19] However, as I will discuss along with the estimation, it seems as though a partisan gerrymanderer need not worry about other demographics after controlling for how "Republican" a district is.

3.1.2   Statistical Model and Estimation

Let $\bigwedge(\cdot)$ be the logistic function. In order to facilitate estimation, assume that $F(x_i) = \bigwedge(\alpha + \beta x_i)$. β is a vector of coefficients and α is some constant. Now, $F(\cdot)$ can be estimated according to the standard logit model.

Prior to discussing the results of the estimation, I should briefly discuss incumbency, a variable conspicuously absent up until now. In general, incumbents nearly always win if they run for reelection. For instance, between 1972 and 1990, there was a Democratic incumbent in 2140 House races. The incumbent won 93.178% of these races.[20] However, all incumbents eventually do not run again (retirement, death, scandal...). Furthermore, districting schemes tend to be long lived (10 years). So, there is a strong possibility that any seat will become an open seat under the implemented scheme. Therefore, gerrymanderers may be particularly concerned with how their party would fair in each district were the seat there open. Then, one should think of $F(\cdot)$ as the probability of electing a Republican, conditional on the seat being open. Since the gerrymanderer may not know who will be retiring, facing a scandal, or dying over the next ten years, they should maximize the sum of the probabilities that they win each seat conditional on that seat being open. After all, the redistricter has far more control over elections to open seats.

Alternatively, the gerrymanderer may be particularly concerned with the next election. In this case incumbency plays an important role and it may be

---

[17] This dataset contains data going further back in time. However, I chose not to use it as Wallace's 1968 run for president and concerns about the earlier Dixiecrats could have confounded my "Republicaness" measure.

[18] This data is available for all of the more recent elections in pdf format in CQ Weekly publications, released around mid-April following each congressional election. Future versions of this work will incorporate this more recent data.

[19] Currently available at http://web.mit.edu/17.251/www/data_page.html#5

[20] 1328 such races featured a Republican incumbent. In these races the incumbent won 92.169% of the time.

Tab. 2: Summary Statistics

	r_i	% Black	% Hispanic	% Urban	Median Age
Mean	-0.00611	0.102	0.0489	0.726	29.274
Std Dev	0.116	0.146	0.0781	0.219	3.207
obs	333	333	283	333	333

(a) Incumbentless

	r_i	% Black	% Hispanic	% Urban	Median Age	RepIncumb	DemIncumb
Mean	-0.00634	0.110	0.060	0.733	29.329	0.343	0.560
Std Dev	0.115	0.144	0.104	0.227	3.140	0.475	0.497
obs	3429	3429	3077	3429	3429	3429	3429

(b) With Incumbents

appropriate to include variables related to incumbency in the $x_i's$.

Therefore, I estimate $F(\cdot)$ under both assumptions. I base my estimates only on elections in the years for which I currently have data on election outcomes (1972-1992). Also, it may have been possible that some districts were designed specifically with some information about the immediately upcoming elections in mind. Thus, in order to remove this possible source of bias, I ignore years ending in 2.[21] I also ignore the few districts which elected independents. This leaves me with 3429 U.S. House races off of which to estimate $F(\cdot)$. The specifications focusing on open seats limit the data to 333 U.S. house races. I also estimate $F(\cdot)$ including demographic variables other than r_i, specifically: percent Black, percent Hispanic, percent urban, and median age. Data on the number of Hispanics was missing for some districts. Therefore, the estimates using the additional demographics were based on fewer U.S. House races. I provide summary statistics for my variables in incumbent-less elections during the relevant years in Table 2.

The results of my estimation are presented in Table 3. Estimated coefficients are presented with standard errors in parentheses beneath each estimated coefficient. In the first two columns I report results for districts without incumbents. In column (1) I present the results based on using r_i as the only demographic of interest. Clearly, districts that are more "Republican" are significantly more likely to elect a Republican to represent them in Congress. Also, it appears that districts which were about as Republican as the nation as a whole tended to

---
[21] Also, my results from the 1980 presidential election by Congressional District would not have matched up with the Congressional Districts in 1982, due to the intervening redistricting.

favor Democratic candidates over this time period (the estimated constant was significantly less than 0).

Additionally, I explore the value of controlling for other demographics in column (2) by including the median age of residents as well as the percentage of residents who were Black, Hispanic or Urban. While Black voters tend to choose Democratic candidates, controlling for this demographic does not seem to improve my estimates. In fact, none of the added coefficients are significant at even the 10% level. Performing a Wald test on the restriction that all of the added coefficients are zero yields a p-value of 0.491. This test further suggests that controlling for other demographics adds little to my specification of $F(\cdot)$. Since I'm already controlling for how "Republican" a district is, also controlling for another variable (such as % Black) which predicts how "Republican" a district is does not lead to better predictions of electoral outcomes. Therefore, for the rest of this paper I will assume redistricters ignore other demographic variables.

The last two columns report results for the estimation when controlling for incumbency. This was done by including dummy variables for the presence of Republican or Democratic incumbents and two variables interacting r_i with those dummies. Column (3) ignores demographics beyond r_i and incumbency. Here, the coefficients on both dummies have the expected sign and are significant at the 1% level. The coefficients on both of the interaction terms are negative and significant. The affect of "Republicaness" on electoral outcomes appears to be muted in the presence of incumbents.

Alternatively, I include additional demographics in column (4). Here one of the added coefficients is significant at the 10% level (with a p-value of 0.098). However running a Wald test on the restriction that all of the added coefficients are zero yields a p-value of 0.337. As such for the rest of this paper I will use columns (1) and (3) as my estimates of $F(\cdot)$.

With $F(\cdot)$ estimated it is now possible to determine the theoretical solution to the optimization problem from equation 2.1 using Theorem 1. Assuming the $F(\cdot)$ implied by column (1) from Table 3.[22] I defined $\underline{r}$ and $\overline{r}$ such that Bush's share of the vote in any district was bounded by $[0, 1]$. Checking all candidate solutions implies that the theoretically optimal Republican districting scheme for Pennsylvania in the 2000's is one in which six districts were completely "Democratic" (Gore would have received 100% of the vote) and $r_i = 0.201$ in

---

[22] This Theorem does not apply to the specification from column (3) as it does not allow for variables like incumbency.

Tab. 3: Estimation Results

	(1)	(2)	(3)	(4)
r_i	13.090***	14.621***	13.090***	14.779***
	(1.790)	(2.256)	(1.790)	(2.212)
$r_i * RepIncumb$			-9.661***	-11.365***
			(2.316)	(2.682)
$r_i * DemIncumb$			-6.594***	-7.565***
			(2.064)	(2.505)
$\%Black$		-1.501		-0.798
		(1.754)		(0.707)
$\%Hispanic$		-0.987		-0.773
		(2.057)		(0.809)
$\%Urban$		-1.002		-0.537*
		(0.750)		(0.324)
$Median\ Age$		0.017		-0.005
		(0.046)		(0.021)
$RepIncumb$			2.672***	2.816***
			(0.182)	(0.205)
$DemIncumb$			-1.995***	-1.830***
			(0.161)	(0.185)
$Constant$	-0.475***	0.436	-0.475***	-0.037
	(0.133)	(1.402)	(0.133)	(0.655)
obs	333	283	3429	3077

*Significant at the 10% level
**Significant at the 5% level
***Significant at the 1% level

the rest of the districts. (Bush would have received about 70% of the vote in these districts.) The estimates from column (1) imply that the Republicans could have expected to win 11.656 of the 19 seats under this plan.

3.2   Generation of Sample Plans

In order to perform my statistical test I generate a large sample of plans randomly drawn so as to respect contiguity and population equality. I then evaluate each plan according to the objective function developed in Section 2 and my estimates of $F(\cdot)$ from Section 3.1.

3.2.1   Data

I make use of two Census 2000 TIGER/Line shapefiles publicly available from Esri.[23] Each shapefile represents the state of Pennsylvania broken down by geographic boundaries from the year 2000. One splits the state into 9418 precincts,[24] while the other breaks the state into 3135 census tracts. Additionally, I merged the precinct shapefile with the 2000 election results and demographic data by precinct available from the Federal Elections Project.[25] (Lublin and Voss, 2001) I also merged the tract shapefile with the 2000 census population counts by tract downloaded through the American Fact Finder available on census.gov. Additionally, I downloaded a shapefiles representing Pennsylvania broken down by 2000 and 2002 Congressional District from the Census Bureau.[26]

Precincts are the smallest level at which votes are counted. This dataset then allows me access to the demographics of interest at the most disaggregated level available to a redistricter. Unfortunately, the algorithm I discuss in Section 3.2.2 took too long to divide 9418 precincts into 19 contiguous, equipopulous districts. Therefore, I used the elections data at the precinct level to estimate the number of votes received by Bush and Gore in each census tract. I estimated the number of votes received by a candidate within a tract as the sum of the votes received in each precinct times the percentage of that precinct within the tract.[27]

---

[23] Currently all available from http://arcdata.esri.com/data/tiger2000/tiger_download.cfm

[24] referred to as "voting districts"

[25] The merge was generally straightforward, but the two datasets didn't always match up exactly. Please contact me for access to this data as well as a readme file explaining my approach to the merge.

[26] http://www.census.gov/geo/www/cob/bdy_files.html

[27] So, if tract 1 contained 50% of precinct A and 25% of precinct B and no other area, then I estimated the number of votes received by Gore in tract 1 as .5 times the number of votes

Looking at the two maps, it appeared that a large majority of all precincts lay entirely within a larger tract. Therefore, this method "estimated" which tract got most precincts' votes with zero error. I also used the same method with the precinct shapefile and the 2002 Congressional District shapefile to estimate the number of votes Bush and Gore received in each 2002 Congressional District in Pennsylvania. Precincts were also unlikely to lie in multiple Congressional Districts. This estimate should also be very accurate by a similar argument.

Additionally, I used the same method to estimate the population of each tract living in each of the 21 Congressional Districts used for the 2000 elections. Generally, a tract was entirely within one district. So, it's population was accurately counted as all living in that district. For those interested in the GIS techniques which led to these estimate please see Appendix B.

3.2.2 Algorithm

The first part of my algorithm follows directly from Cirincione, Darling and O'Rourke (2000). It starts with the map of Pennsylvania by census tract with no tracts assigned to any Congressional District. First, it randomly chooses an unassigned tract and adds it to District 1. Then it randomly chooses an unassigned tract which neighbors the now growing District 1 and adds it to the district. This last step repeats until the population of the district reaches the ideal district population (646,371), or there are no more unassigned neighbors.[28] Then the above repeats for each of the other 19 districts. (It randomly selects an unassigned tract and assigns it to district 2 and then randomly selects an unassigned neighbor...). If any tracts are left unassigned at the end they are randomly chosen and assigned to a neighboring district.[29] The original Cirincione, Darling and O'Rourke (2000) algorithm required that the population of each district be within 1% of ideal or the algorithm would simply restart, throwing out the plan and starting from scratch. Under this restriction my algorithm always started over. Instead, my algorithm only restarts if at least one district isn't within 50% of the ideal district population. Otherwise, it attempts to fix

---

received in precinct A plus .25 times the number of votes received in precinct B.

[28] There may be no unassigned neighbors if the construction of an earlier district caused there to be a few unassigned tracts surrounded by already assigned tracts. If the algorithm started creating a district from one of these tracts, it would not be able to make it large enough before running out of unassigned neighbors.

[29] It is possible that all of the tracts neighboring an unassigned tract are also unassigned. In this case the algorithm moves on, randomly selecting other unassigned tracts until it finds one with a neighbor assigned to a Congressional District. If the tract chosen has multiple neighbors assigned to different districts, the district is chosen randomly from among the possible choices.

any population inequality by randomly choosing tracts which neighbor another district and reassigning them if the neighbor has a lower population. (If a tract neighbors multiple other districts, one of the neighbors is randomly chosen.) This continues until the population of all Congressional Districts is within 5% of the ideal district size.

One can think of plans generated by this algorithm as samples of the plans that might be produced if we tasked an otherwise disinterested cartographer with dividing a state into equipopulous, contiguous districts. She might start picking a small part of the state and decide that should be in one district and then subsequently add small neighboring bits of the state until the district was large enough. Then she might repeat the same process for the other districts she was required to create. If she got stuck, instead of starting from scratch she might choose to reallocate small areas until she achieved an appropriate solution.

3.3   Results

My approach generates districting schemes which are apolitical and random, but which also respect the population equality and contiguity constraints required of legal plans. These are exactly the types of alternative plans against which one should compare the actual plan. The purpose of the test is to determine whether partisan considerations guided the design of the redistricting scheme. Again, the null hypothesis is that partisan considerations were not used to develop the plan for Pennsylvania. If the chances of a disinterested cartographer producing a scheme so favorable to the redistricting party were less than 5%, then we should reject the null hypothesis at the 5% level. I would then be forced to accept the alternative, that partisan considerations were used to develop the plan for Pennsylvania. This would imply that charges of partisan gerrymandering are valid by definition.

The program to implement my algorithm was coded in R and my code made use of packages designed to deal with this type of geographic problem.[30] Using my program I generated a sample of 10,000 equipopulous, contiguous districting schemes for Pennsylvania.[31] For each sample plan, I used the list of tracts contained within each district and the estimated election results by tract

---

[30] Specifically, see Altman and McDonald (2011)

[31] This took about 2 days. As of the time I'm writing this, my program is currently re-running to generate a new sample of 10,000 districting schemes. It seems unlikely that the change will affect my results, but the larger sample will be used in future versions of this work.

Tab. 4: Results

Specification from Table 3	(1)	(3)
Range of Expected	6.061 to 9.498	6.920 to 12.282
Actual Expected	7.938	10.707
Percentile of Actual	76.62	96.25

to calculate the number of votes received by Bush and Gore in each district.[32] From there, the calculation of r_i for each district in each sample plan (and the actual plan) was straightforward.

I used the estimates of how many people in each tract lived in each of the pre-redistricting districts similarly to determine the population of each new district that came from each of the old districts. I assumed that incumbents would then run in whichever new district contained the largest number of their former constituents. From here it was straightforward to determine which districts had incumbents from which party.[33] Incumbency dummies were generated for the actual plan by looking at which incumbents ran in which of the new districts.[34]

Then, using the estimates of $F(\cdot)$ from columns (1) and (3) of Table 3, I estimated the value of the partisan redistricter's objective function, Equation 2.1, for each sample plan as well as for the actual 2002 Congressional Districting scheme.

The results are reported in Table 4. When ignoring incumbency (Specification (1)), the estimated value of the objective function among the sample plans ranged from 6.061 to 9.498. The same value for the actual plan was 7.938. This was at the approximately the 77th percentile of the value for the sample plans. Therefore, using this specification I cannot reject the null hypothesis that partisan considerations were not used to develop the plan for Pennsylvania. If one thinks redistricters are farsighted and we should not be considering the interaction between redistricting and current incumbents, then my test cannot prove intent. However, it still shows that the chosen plan was particularly favorable for Republicans. These results suggest that only a small fraction ($\approx$23%) of

---

[32] Technically, these numbers are estimates as the number of votes in each tract are estimates. However, as discussed earlier those estimates generally (though not always) should have zero error.

[33] Sometimes, two Democratic (Republican) incumbents would have decided to run in the same district. In this case the *DemIncumb* (*RepIncumb*) dummy equaled 1.

[34] Of the 21 Representatives elected in 2000, there was only one that retired without running in 2002. However, most of his constituents were combined with the district of another Democrat (likely prompting the retirement). So, my method would have produced the appropriate incumbency variables under the actual plan as well.

possible plans were better for Republicans than the one they chose.

The results are much stronger under Specification (3). Republicans could expect to elect more Representatives under their chosen plan than under 96.25% of the sampled plans. We should then reject the non-partisan null-hypothesis at the five percent level. We must then accept the alternative, that partisan considerations were used to choose the redistricting plan. This implies that the plan is a partisan gerrymander by definition. Bear in mind that this is a working paper and these results should be considered preliminary.

4 Discussion

Assuming gerrymanderers are only concerned with the next election, my test confirms that the Pennsylvania plan was in fact a partisan gerrymander. However, if gerrymanderers are concerned with the long run, it may make sense for them to forgo considering the current incumbents while designing their plan. They should then be concerned with how candidates from their party would do in an election to an open seat (which all seats may eventually become).

However, my results suggest that the Republicans redistricting Pennsylvania in 2001 were especially concerned with the next election. This seems to match with Pennsylvania's electoral history for the rest of the decade. While in 2002 and 2004, Republicans captured 12 of Pennsylvania's 19 seats, this trend was reversed in the next two elections. Under the same districting scheme in 2006 the Democrats captured 11 seats, and then 12 in 2008. The chosen plan worked well in the 2002 election, but eventually Democrats were able to elect a majority of the state's Representatives.

This suggests that gerrymanderers are particularly interested in the next election. In Pennsylvania it seems Republicans chose the plan under which they would do the best in the immediately following election. If true, then the appropriate specification is likely the one under which I was able to confirm partisan gerrymandering.

Also, in other cases even if my test fails to show intent, it still provides valuable information about effect. If courts were to decide that plans which significantly favored one party were illegal because they had the effect of a partisan gerrymander they could easily use a less restricted version of my test. An obvious criteria to use might be to require that any plan yield a test statistic within the middle 50% of those of sampled plans. This criteria would not be based on proving intent, but on demonstrating effect. This rule would ban

plans under which the redistricting party did better than could be expected, "on average," under plans generated by apolitical processes. While this would not eliminate the possibility of partisan gerrymandering, it would add an additional constraint. For instance, the plan for Pennsylvania tested here would fail to meet this stricter criteria under both specifications.

While my test makes some rather specific assumptions, the method is easily adapted to alternative conceptions of partisan gerrymandering. The theory section could be readily modified to use another objective function.[35] For instance, one might want to include some level of risk aversion, where the redistricter is particularly concerned about how well they do when the opposing party has a strong election year. Empirically, one might want to specify a different functional form for $F\left(\cdot\right)$, or include more demographics in the estimation. In fact, earlier in this project I estimated $F\left(\cdot\right)$ using kernel density procedures, avoiding any functional form assumptions at all. However, with very little data off of which to base estimates in the the tails of the distribution of r_i, it seemed wise to switch to parametric methods. This may be less of a concern once I update the dataset to include more recent elections. Also, alternative random algorithms could be used to generate the sample of alternative plans. Perhaps one thinks a disinterested cartographer would approach the problem differently than the algorithm proposed here.

4.1   Limitations

Again, all results should be considered preliminary. This paper is a work in progress. Beyond that, several potential criticisms of my approach seem apparent. Obviously, opinions may differ as to the actual objective of a partisan gerrymanderer. I believe I used the most obvious choice: maximizing the expect number of Representatives from the gerrymanderer's party. However, nothing about my test requires this specific objective function, and an alternative could easily be used if preferred.

Perhaps the most significant concern with my approach is that the redistricter could have access to some unobserved information about the population or the upcoming elections. I deal with this as a source of bias for my estimation by only estimating off of elections at least one cycle removed from redistricting. It seems unlikely that in 2001 redistricters had special information about the elections that were not going to occur until 2004.

---

[35] So long as that function could be estimated using available data.

It is still possible that redistricters had some unobserved information that could have made a particular plan more (or less) attractive than my estimates would suggest. While this is a potential source of error for my test, I argue one should only expect it to increase the incidence of type II errors. If redistricters were not using partisan objectives to choose a plan, there is no reason to suspect that any additional information they had would tend to lead them to choose a plan that looked especially partisan based on my estimates (thus causing a type I error). The alternative is that the redistricter was using partisan objectives to choose a plan. If this was the case and the unobserved information caused an error, it would have to be a type II error (failing to reject non-partisan motives when the redistricter was in fact partisan).

Currently the estimation of $F(\cdot)$ is based on only one demographic and less than current data. In Section 3.1.2 I made the case that one should treat deviation from national Republican presidential vote share as a sufficient statistic for other demographics that may be of interest to a partisan redistricter. However, there is nothing to prevent one from using more information in the estimation of $F(\cdot)$.

I assumed an incumbent would run in a newly created district if that district contained more of his former constituents than any other district. This may be an inappropriate assumption. Suppose 51% of an incumbents constituents ended up in district 1 and 49% ended up in district 2. My assumption implies that incumbent will run in district 1 during the next election. That implication may be wrong, especially if district 2 would be an otherwise open seat whose demographics favor that incumbent's party. I could remedy this issue by formally modeling the incumbent's choice of which district to run in. I hope to do so in future versions of this works. However, in the case of Pennsylvania, this would be a game with 21 players, each choosing 1 of 19 districts in which to run. Accounting for this feature may unnecessarily complicate the model.

Ideally, statistical tests are based on random samples drawn from some population. There is nothing to guarantee that my algorithm generates plans which are randomly drawn from the set of contiguous equipopulous plans. Given the size of the redistricting problem, generating such a random sample seems computationally infeasible. As already discussed, there is no feasible way to generate the entire population of such plans, and taking a random sample would seemingly involve randomly assigning each tract to a Congressional District and then throwing out the plan and starting over if you had not created a contiguous, equipopulous plan. Such an approach would almost always create an invalid

plan and have to start again, essentially ad infinitum.

Instead, I randomly generate a sample of plans using an algorithm which I believe mimics how a disinterested cartographer might go about creating legally valid districts. There may be some disagreement as to how a disinterested party would in fact attempt to draw a legally valid plan. Nothing about my test relies on this specific algorithm; it could be substituted for another that seemed more appropriate. Additionally, it might be prudent to run my test using multiple algorithms, as a robustness check to ensure my results aren't driven by a particular choice of algorithm.

Specifically, the current algorithm starts with a blank map to randomly draw districts. It may be more appropriate to assume the *re*districter would start with the previous map and modify it to satisfy the current constraints. This difference may partially explain why my results are so surprisingly strong when accounting for incumbency. There were 11 Republican incumbents going into the 2002 House elections in Pennsylvania. The actual plan did not "waste" any Republican incumbents; each was placed in a separate district. This outcome seems unlikely when redrawing the map from scratch. However, it seems more likely if the new map were drawn by making modifications to the old map until it satisfied the new requirements. In future versions of this paper I plan to also use such an algorithm as a robustness check.

4.2   Further Work

The states are currently going through another round of redistricting based on the results from the 2010 census. In several states this process is completely controlled by one party. Hopefully, this work can be helpful in evaluating claims of partisan gerrymandering based on these plans. Ideally, the Supreme Court might find a test based on this one acceptable for evaluating charges of partisan gerrymandering.

As mentioned earlier, scholars have already used random redistricting to attempt to evaluate charges of racial gerrymandering, but they chose their test statistics arbitrarily. My methodology could be readily applied to improve this line of literature. The most obvious objective function that we might ascribe to a racially motivated gerrymanderer would be to minimize (or maximize) the number of minority representatives. Clearly, one would need to use different demographic variables when estimating this objective function, but the basic approach would remain the same.

Additionally, there are a few tasks that, once completed, will improve this paper. I need to update my election results data so that I may base my estimates off of more recent data. Though it is not clear why, it is possible that the relationship between preference for Republican presidential candidates and Republican congressional candidates may have changed since the 1970's. Also, running the test based on samples generated by multiple algorithms would provide an important robustness check.

5   Conclusion

I developed a statistical test for partisan gerrymandering and applied it to the districting plan chosen by Pennsylvania in 2001. My approach remedies a universal problem among papers which use random redistricting to evaluate gerrymandering claims: arbitrary test statistics. I treat partisan gerrymandering as a maximization problem faced by a seat maximizing party. I estimate the parameters of this objective function, and my test rests naturally on the value of the objective function under the actual and simulated plans. Using GIS techniques, I randomly produce a large sample of legally valid districting plans and calculate the value of the necessary demographics to evaluate the objective function under each sample plan.

When controlling for incumbency, my test confirms that the plan chosen by Pennsylvania in 2001 was a Republican partisan gerrymander. Since Vieth v. Jubelirer (2004) courts have been unable to hear claims of partisan gerrymandering for lack of an acceptable test. This test should serve their purposes well. It avoids many of the messy questions related to how districts "should" be drawn and what plans would be "fair." Instead, it focuses on how unlikely it would be to see a plan so favorable to the redistricting party developed by a disinterested cartographer.

References

Altman, M., and M.P. McDonald. 2011. "BARD: Better Automated Redistricting." *Journal of Statistical Software. Forthcoming, URL http://www. jstatsoft. org*, 42(4): 1–28.

Cirincione, Carmen, Thomas A. Darling, and Timothy G. O'Rourke. 2000. "Assessing South Carolina's 1990s congressional districting." *Political Geography*, 19(2): 189–211.

Coate, Stephen, and Brian Knight. 2007. "Socially Optimal Districting: A Theoretical and Empirical Exploration*." *Quarterly Journal of Economics*, 122(4): 1409–1471.

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): 812–841.

Dopp, Kathy Anne. 2011. "Legislative Redistricting - Area and Population Compactness and Population Density Distribution Measures." *SSRN eLibrary*.

Engstrom, Richard L., and John K. Wildgen. 1977. "Pruning Thorns from the Thicket: An Empirical Test of the Existence of Racial Gerrymandering." *Legislative Studies Quarterly*, 2(4): 465–479.

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

Garfinkel, R. S., and G. L. Nemhauser. 1970. "Optimal Political Districting by Implicit Enumeration Techniques." *Management Science*, 16(8): B495–B508.

Gelman, Andrew, and Gary King. 1994. "Enhancing Democracy Through Legislative Redistricting." *The American Political Science Review*, 88(3): 541–559.

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

King, Gary. 2006. "Elections to the United States House of Representatives, 1898-1992."

King, Gary, and Robert X Browning. 1987. "Democratic Representation and Partisan Bias in Congressional Elections." *The American Political Science Review*, 81(4): 1251–1273.

Lublin, David, and D. Stephen Voss. 2001. "Federal Elections Project."

Rossiter, D J, and R J Johnston. 1981. "Program GROUP: the identification of all possible solutions to a constituency-delimitation problem." *Environment and Planning A*, 13(2): 231–238.

A Proof of Theorem 1

Since the theorem considers a continuous objective function evaluated over a compact subset of R^N, we are assured that at least one solution exists. I prove this theorem by contradiction. In the first two of the following cases I show that no solution could satisfy particular properties which would not satisfy the theorem. Then, in Case 3 I show that the only other class of potential solutions which could violate the theorem either admits a solution which does not violate the theorem, or cannot be optimal. Let z be the value of the demographic about which $F(\cdot)$ is symmetric.[36]

Case 1. First suppose that there is a solution $\{r_i\}_{i=1}^N$ such that for some i and j:

$$r_i \neq r_j \quad and \quad r_k \in [z, \overline{r}] \ \forall k \in \{i, j\}.$$

WLOG assume $r_i < r_j$. However note that $F'(r_i) > F'(r_j)$ by assumption. Therefore, it would be possible to decrease r_j by ϵ and increase r_i by ϵ while holding the demographics of the other districts constant and increasing the overall value of the objective function. This is a contradiction. Therefore, in any solution, all districts whose demographic is at least z have the same value for the demographic.

Case 2. Similarly, suppose that there is a solution $\{r_i\}_{i=1}^N$ such that for some $i \neq j$:

$$r_k \in (\underline{r}, z) \ \forall k \in \{i, j\}.$$

WLOG assume $r_i \leq r_j$. However note that $F''(v) > 0 \ \forall v < z$ by assumption. Therefore, it would be possible to increase r_j by ϵ and lower r_i by ϵ while holding the demographics of the other districts constant and increasing the overall value of the objective function. This is a contradiction. Therefore, no more than one district can have a demographic strictly less than z, but strictly greater than $\underline{r}$.

Case 3. Suppose that there is a solution $\{r_i\}_{i=1}^N$ which does not exhibit Cases 1 or 2. Also, suppose that for some $i \neq j$:

$$r_i \in (\underline{r}, z) \quad and \quad r_j \in [z, \overline{r}).$$

If $z - r_i \neq r_j - z$ then note that $F'(r_i) \neq F'(r_j)$ by assumption. Therefore,

---

[36] i.e. $\forall y \geq 0 \ 1 - F(z + y) = F(z - y)$.

it would be possible to increase r_j by ϵ and lower r_i by ϵ^{37} while holding the demographics of the other districts constant and increasing the overall value of the objective function. This is a contradiction. Alternatively, if $z - r_i = r_j - z$, then if we were to change r_i and r_j to equal z while holding the demographics of the other districts constant we would not change the value of the objective function and the constraint would still hold (both by the symmetry of $F(\cdot)$ about z). Either there exist other districts in this new solution with demographics strictly greater than z, or there do not. In the former we would have a contradiction by the logic of Case 1, in the latter we would have an alternative solution that satisfies Theorem 1

By Cases 2 and 1, in any solution there can only be one value of the demographic weakly greater than z and all but one district whose demographic is strictly less than z must be at $\underline{r}$. Therefore, any solution must satisfy Theorem 1 or exhibit Case 3. As any solution exhibiting Case 3 must admit a solution which satisfies Theorem 1, there must always exist a solution that satisfies Theorem 1.

B   Estimating votes in other Shapefiles

In order to get my estimate of the votes Bush and Gore received in each census tract I did the following in ArcGIS: First, I added an area attribute to the precinct shapefile which was a measure of the area of each polygon. Next, I did a union with tract shapefile. This split up any precincts into the parts strictly within a tract. Each polygon in the new shapefile then had attributes recording the number of votes in, and area of, the precinct from which they came. Next I added an area attribute to the union shapefile which measured the area of each new polygon (using the same units as the area measure for the precincts). Next, I generated new attributes to estimate Gore and Bush votes in the each polygon in the union shapefile. Specifically, estimated votes received in a polygon in the union were set equal to the votes received in the precinct the polygon came from times the ratio of the area of the new polygon to the area of the original precinct. The polygon union shapefile was then converted to a point shapefile (all attributes were assigned to a point at the centroid of each polygon). Finally, I used the join by location feature to sum the estimated votes from each point that occurred within a polygon in the census tract shapefile. The analogous approach was used to estimate votes by 2002 Congressional District.

---

[37] Depending on the direction of the inequality between $F'(r_i)$ and $F'(r_j)$, ϵ will need to be either positive or negative.