All Equal Sized Single-Member Districts are not Created Equal:

An Analysis of Electoral Skew in the U.S. House

Michael C. Dougal*

Abstract

Even when population is uniformly distributed across districts, preferences are not. The U.S. House contains more districts that are extremely liberal than extremely conservative, but has more slightly conservative than slightly liberal districts. Using Democratic presidential vote share as a proxy for district preferences, this paper estimates the partisan impact of the skewed distribution of electoral preferences across districts by comparing actual election outcomes to simulated election results in which districts are randomly drawn from a normal distribution with the same mean and standard deviation of preferences across districts as the actual House districts of that year. A key finding of this paper is that the distribution of preferences across districts has been consistently skewed in favor of the Republican Party and that when the Republican Party has controlled the Speaker's gavel the magnitude of this effect has typically been larger than the Republican margin of victory.

Prepared for presentation at the Annual Meeting of the Western Political Science Association, Hollywood, CA, March 28-30, 2013.

* Graduate Student, Travers Department of Political Science, University of California, Berkeley, mdougal@berkeley.edu

"How Democratic Is the American Constitution?" Not very democratic according to Robert Dahl, but his answer has little to do with House elections (Dahl 2003). That is because much of the normative and empirical work in the academic literature on what makes a fair legislative election and how to measure it focuses on the extent to which the distribution of seats to districts fairly reflects the population distribution between districts. By this measure, the U.S. House of Representatives is a well-apportioned legislative body in which every citizen's vote is counted almost equally. Certainly in contrast to the U.S. Senate, which Samuels and Snyder (2001) find to be the 5th most malapportioned upper chamber in a comparative study of 78 countries, the House's close to equal sized districts by population appear fair and democratic. Population imbalances are not the only way, however, for representation to become skewed. Even if population is uniformly distributed between districts, preferences are not.[2] If like minded supporters of one major party are disproportionately clumped in a small number of districts, while the opposing party's supporters are more evenly distributed across all districts, the former party is at a disadvantage in two-party single-member plurality districts because many of their supporters' votes will be wasted in one-sided districts, while the latter party will win more seats by smaller margins, thus wasting fewer votes.[3] All equal sized single-member districts are not created equal.

In the case of the United States, voters with more liberal preferences are more heavily concentrated in particular house districts, while conservatives are more evenly

[2] It preferences were uniformly distributed between single member districts the results would be disastrous for the minority party, which would fail to win a single seat no matter how close it was to becoming the majority party.

[3] Unless otherwise noted, parties discussed in this paper are one of only two major parties and all districts are single member plurality districts in which one of the two major parties' candidate wins. Vote share thus refers to two-party vote share unless otherwise specified.

distributed across all districts (Rodden and Warshaw 2009). The median house district has voted more conservatively than the nation as a whole in every presidential election since 1952 (Mayhew 2011). Mayhew argues that the impact of this skew, which on average is a 1.1 percent difference between the two party presidential vote share in the median house district and the nation as a whole, has had a relatively small impact on American politics. Rodden and Warshaw (2009) judge this skew as more problematic, but hypothesize that it has benefits for the Democratic Party as well as the Republican Party. Mainly, they suggest that "it would appear that there is an asymmetry whereby the Democrats make bigger gains in conservative districts during their moments of strength than do Republicans in liberal districts, and even in their time of greatest strength, the Republicans did not extend their reach very far across the median" (Rodden and Warshaw 2009: 36). While Mayhew examines the similarity of House electoral outcomes to Senate electoral outcomes and Rodden and Warshaw provide an excellent description of the distribution of voter preferences across districts, neither focuses on offering a counterfactual, in which there is no skew, against which to measure the partisan impact of electoral skew.

The aim of this paper is to empirically test the impact of the skewed distribution of preferences across congressional districts on party size in and party control of the U.S. House of Representatives by providing an appropriate counterexample in which preferences are not skewed. Building on methods developed by Gudgin and Taylor (1980) in the British context, I compare the electoral outcomes of actual House elections between 1952 and 2012[4] to hypothetical House elections with the same mean and

[4] The data on House elections and two-party presidential vote share by congressional district used in this paper were generously shared with me by Gary Jacobson.

standard deviation of Democratic presidential vote share across districts, but normally distributed. This method captures the impact of the non-normality of the distribution of preferences across districts, often referred to in the literature as distributional bias or electoral skew.[5] A key finding of this paper is that the distribution of preferences across districts has been consistently skewed in favor of the Republican Party and that when the Republican Party has controlled the Speaker's gavel the magnitude of this effect has typically been larger than the Republican margin of victory.

2) Literature Review

The literature on electoral bias often uses some sort of votes to seats ratio at the aggregate level to determine the extent to which electoral outcomes are skewed. If the proportion of votes is equal to the proportion of seats, the election is considered unbiased, while a party that picks up more seats per votes than the opposing party benefits from an electoral bias. The problem with any method that uncritically compares the number of votes to the number of seats obtained is that the relationship between votes and seats is not linear, but rather is an s-shaped curve that will consistently deliver more seats per vote to the majority party (Grofman 1983). Sometimes called the "winner's bias," this relationship is well known (Gudgin and Taylor 1980: 518). It is often said to be a product of the "cube law," which "says that the ratio of assembly seats of two major parties is approximately the cube of the ratio of votes" in single member district parliamentary elections (Taggepera 2004: 257). The cube law has been found to be a relatively good fit for the British case in which Kendall and Stuart (1950) first introduced it to the academic

[5] In this paper, electoral skew specifically refers to the non-normality of the distribution of preferences across districts.

literature, but Tufte (1973) showed that in most cases a logistic curve is actually a better fit. At any rate, it is widely agreed that the marginal effect of a one point increase in vote share on seat share is much greater around 50 percent and of diminishing marginal utility to the majority party as they come closer and closer to 100 percent of the vote.

Due to the s-shaped curve of a seat to vote ratio, the actual ratio alone will only get us so far. While Tufte (1973: 547) defines an unbiased electoral system as one in which the marginal utility to an increase in seat share to vote share does not change and is thus a one-to-one ratio between votes and seats, this definition of bias tells us only that single-member plurality districts violate the one person one vote principle in a way that a single national multi-member proportional representation district would not. In a majoritarian election system with single-member districts, however, there is a winner's bias. Barring a switch to proportional representation, the question is not is there a winner's bias, but does each party benefit from the winner's bias in equal measure? To answer this question, Brookes (1960) introduced a measure of electoral bias that did not examine the raw seats-to-votes ratio, but instead asked which party had a better ratio if they each received the same percentage of the vote. Of course, two parties are exceedingly unlikely to receive the same number of votes in an election, so in order to construct a counterfactual in which the parties did obtain an equal number of votes, Brookes applied a uniform shift across districts in order to obtain comparable cases. For example, if the Labour Party received 57% of the two-party vote by district, Brookes would either subtract 7% from every district to examine a 50/50 election or subtract 14% from every district to examine an election in which the Conservative Party received the same 57% as Labour did in the actual election.

Brookes' method has been applied in the American case by Johnson et al. (1999), but its application was highly problematic. Brookes' method was developed in the context of a Westminster-style parliamentary democracy like that of the United Kingdom or New Zealand, not a presidential system like the United States. The important difference for Brookes' method is that all legislative districts are contested when control of the legislative chamber also determines control of the executive government, but many seats in the U.S. House go uncontested because there is no broader incentive for a party to contest a House election if that district is highly skewed towards the other party.[6] This presents a missing data problem for calculating the average two-party vote share by congressional district. Johnson et al. (1999) ignore the problem, drop uncontested districts from their analysis, and thus only include 366, 360, and 354 districts in their analysis of House elections in 1984, 1986, and 1988 respectively. In other words, they dropped over 15 percent of the House each year with no discussion of whom they were dropping or why it might be problematic. In each year examined, over 75 percent of the uncontested districts they dropped were Democratic districts. Thus, their conclusions are based on what amounts to a different legislative body than the U.S. House of Representatives.

Missing vote share data from uncontested congressional districts can be imputed when examining seat to vote ratios, though this does require an additional set of assumptions (Gelman and King 1994). However, Brooks' method suffers from a more fundamental problem. As Johnston et al. (1999: 375) acknowledge when using "Brookes'

[6] In the 2010 UK parliamentary elections, there was not a single seat in the House of Commons that went uncontested by either Labour or the Conservatives. Uncontested races at the national level appear to be a peculiarly American phenomenon.

method, the major disadvantage is its assumption of a uniform shift in votes between two parties across all constituencies." While this method will capture the extent to which one party has a set of districts grouped just above or just below the 50% margin of victory, it ignores that electoral skew is in part a function of extreme districts. Applying a uniform shift to all districts ignores that increased support is of little value in an already lopsided district, but that were a party's supporters from a lopsided district in a more even district they would be of much greater marginal utility to the party. Gelman and King (1994) develop a more complex method for estimating partisan bias that includes a random error component such that it does not assume a uniform shift, but to simulate counterfactual election results with the desired vote share they must still model an aggregate partisan swing that shifts the mean vote share to the desired counterfactual. While this shift is not uniformly applied, the underlying model is linear and thus the shift is just as likely to be applied at the extremes as at the center of the distribution, resulting in effectively the same problem.

To provide a counterfactual to an actual election, assumptions must be made. Nevertheless, the assumptions made by either Brookes (1960) or by Gelman and King (1994) change the mean while ignoring that additional support in lopsided districts is of little value to a political party. In contrast, Gudgin and Taylor (1980) propose comparing the seat proportion of a party based on the actual distribution of two-party vote share across districts to a counterfactual scenario in which the electoral districts had the same mean and standard deviation, but were normally distributed. The advantage of this method is that it does not assume a uniform shift or change the variance, but asks how the results would change if district level vote shares of both parties were normally

distributed. The assumption of Gudgin and Taylor's method then is that a normal distribution can be characterized as fair and not partisanly biased. If we were considering minor parties, then this assumption would be unfounded because minor parties benefit from a concentration of supporters and suffer if their supporters are evenly scattered across districts (Gudgin and Taylor 1974). This paper, however is confined to looking at two major parties of relatively equal strength, which do not benefit from over-concentrated support, but rather do better with a more normal distribution of supporters across districts. Given that two parties of equal strength, both normally distributed across districts, would each capture 50 percent of the seats, this is a partisanly neutral measure of the impact of electoral skew. Like both Brookes (1960) and Gelman and King (1994), fairness is thus defined by symmetry. Unlike these other authors, however, Gudgin and Taylor (1980) assume not that they have come up with the correct functional form to change the mean to obtain new distributions with partisanly symmetric means that provide counterfactuals to one another, but that the symmetry of a normal distribution is a fair counterfactual to the actual distribution.

3) Methods

While Gudgin and Taylor's (1980) proposed method accurately measures partisan skew, or more precisely the impact of the non-normality of the actual distribution, it suffers from the same problem as the other methods in the American case, which is the significant number of missing values presented by uncontested seats. To avoid this problem, I propose indirectly implementing Gudgin and Taylor's method by using two party presidential vote share in house districts to predict house election outcomes.

Because the presidency is a nationwide office and the same two major party candidates are on the ballot in every district, Democratic presidential vote share can thus be used as a proxy for a district's ideological preferences, as is standard in the congressional literature. Gudgin and Taylor's proposed counterfactual of a normal distribution with the same mean and standard deviation by district as the actual results can thus be applied at the level of preferences, i.e. presidential vote share, and then used to predict the alternative electoral outcome that Gudgin and Taylor compare to the actual electoral outcome to find the impact of electoral skew. This method has the theoretical advantage that any skew in vote share across districts is in fact a product of the skewed distribution of voter preferences across districts and the practical advantage of relying only on data for which there are no missing values.

Given Mayhew's (2011) finding that the Democratic two party presidential vote share is greater in the nation as a whole than in the median congressional district in every presidential election since 1952, I expect to find democratic presidential vote share to be positively skewed resulting in the Democratic Party consistently winning fewer seats than it would if preferences were normally distributed across house districts. Mayhew argues that electoral skew has not had a profound impact on House elections and the level of skew is relatively minor. Implicitly, he treats a one point difference between the median district and the national mean as if it would change the outcome in roughly one percent of districts, i.e. the districts between the median and the mean. If preferences were uniformly distributed across districts, this would make sense because a one point shift would move one percent of seats across the 50% percent threshold for victory. However, if preferences are more normally distributed, far more than one percent of seats would be

moved passed the 50% cut point for victory by a one percent shift because far more than one percent of cases are centered on the cut point in a normal distribution.

Second, for the number of seats that change hands due to skew to be as small as Mayhew's nominal measure, a one point increase in Democratic preferences, as measured by presidential vote share, cannot result in a more than one point increase in the probability that the Democratic candidate will win. In other words, preferences must deterministically dictate vote share and a seat will only change hands if a uniform shift in preferences pushes it past the 50% + 1 threshold for victory. As Gelman and King (1994) argue, the assumption that preferences deterministically dictate election outcomes is problematic. If, in contrast, districts have a mean preference, but the manifestation of that preference as House vote share in any given election varies slightly according to a normal distribution, then a one point increase in a district's true preference will result in a more than one point increase in their probability of victory around 50%, but will provide diminishing marginal return as the districts mean preferences move farther and farther away from the cut point (this data generating process has the curve of a probit function).

Mayhew's work clearly establishes that the distribution of preferences across districts has been consistently skewed in favor of the Republican Party. While other factors may favor the Democratic Party, such as winning districts with lower average turnout, Mayhew establishes that the non-normality of the distribution of preferences clearly favors the Republican Party. Yet, he does not view this skew as having a large impact, pointing out (prior to the 2010 election) that every time partisan control of the House changed between 1952-2008, control of the Senate changed as well (Mayhew 2011: 29). I predict that a far more significant number of seats is changing hands due to

electoral skew because (1) preferences are unimodally distributed such that more than one percent of the seats are centered at the cut point at 50 percent and (2) I expect that the relationship between the presidential vote share and the party winning at the congressional level will be an s-shaped curve in which a one point increase in the Democratic presidential vote share around the cut point will correspond with a more than one point increase in the probability that the Democratic House candidate gets 50% + 1 in the district.

The rest of this paper can be divided as follows. First, I examine the distribution of democratic presidential vote share and the general shape of its relationship with house electoral outcomes. Next, I use a non-parametric running mean smoother of the actual results by year to assign each district a probability that, given the two party presidential vote share in that district, it will elect a Democrat. I then compare the actual number of seats won by each party under the current distribution of Democratic presidential vote share to a distribution with the same mean (averaged by district) and standard deviation, but which is normally distributed. If the difference between the two is larger than the percentage of seats required to flip control of the House, this measure suggests that control of the House is potentially being impacted by electoral skew. Finally, I repeat this process using national presidential popular vote share, as opposed to the mean presidential vote share across districts used to measure partisan skew, to test whether the Speaker's gavel would change hands under a set up more conservative assumptions that are more biased against such a conclusion.

4) Results

In predicting whether or not a Democrat would win a House seat in a particular district, 50 percent of the two party presidential vote share meant something very different in 1964 than in 1972. While unnecessary when different years are treated separately, pooling two-party presidential vote across different presidential years requires a slight adjustment. Therefore, I "demedian" presidential vote share by year. That is, for each district I subtract the presidential vote share in the median congressional district for that year and add .5 to make results more intuitive to understand. So, for all years a demedianed presidential vote share of .5 is a congressional district with the median presidential vote share for that year.

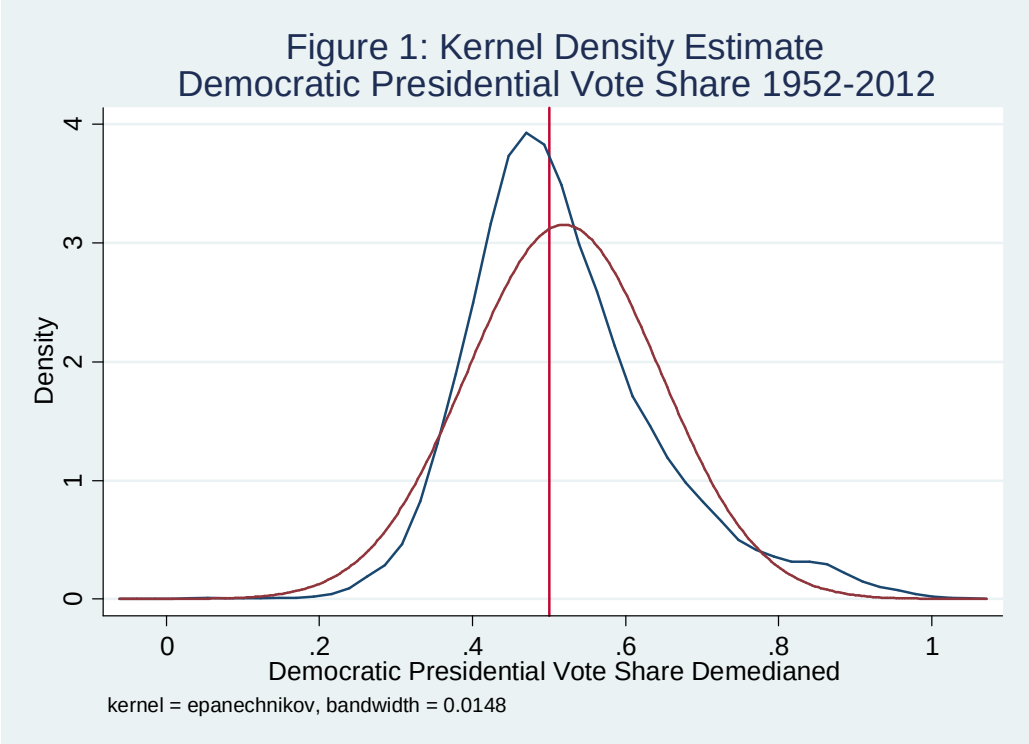

In Figure 1, the blue distribution is the kernel density estimate of the demedianed Democratic presidential vote share, while the red one is a normal distribution with the same mean and standard deviation. The vertical red line at .5 is a reference line at the

sample median. As you can see, Democratic presidential vote share is positively skewed. The kernel density estimate peaks before either the median or mean and there is greater density in the right tail, which represents a disproportionate number of extremely liberal Democratic districts. Indeed, in 2012 there were 73 extremely liberal house districts (Democratic presidential vote share 20% or more above the median house district), but only 20 extremely conservative districts (see Table 1).

Table 1. Districts with Party's Presidential Vote Share 20%+ than Median District

Year	*Democratic*	*Republican*
1952	38	1
1954	38	1
1956	39	3
1958	39	3
1960	32	2
1962	10	0
1964	18	9
1966	10	4
1968	34	1
1970	35	1
1972	29	1
1974	30	1
1976	26	3
1978	26	3
1980	30	7
1982	23	9
1984	33	2
1986	33	2
1988	37	1
1990	37	1
1992	40	9
1994	40	9
1996	41	11
1998	41	11
2000	53	11
2002	54	9
2004	61	9
2006	60	8
2008	57	17
2010	58	19
2012	73	20

Finally, the density estimate is also leptokurtic, or taller than the normal distribution. Thus, there is a larger number of districts with slightly Republican preferences than with slightly Democratic preferences.

While it is clear from Figure 1 that the Democratic Party would do at least somewhat better if preferences were normally distributed across districts, how much better depends on the relationship between presidential vote share and house elections. Figure 2 uses a lowess smoother to graph the relationship between two party presidential vote share demedianed and the Democratic Party winning in a house seat in a particular district. The lowess line reflects the probability that a Democrat will win the seat given the Democratic presidential vote share at that point. The green reference line is at a 45 degree angle and thus is what a one for one relationship between the two variables would look like. The blue dots at 0 and 1 reflect the actual outcomes.

Figure 2: Lowess Smoother
Democratic Presidential Vote Share and House Elections

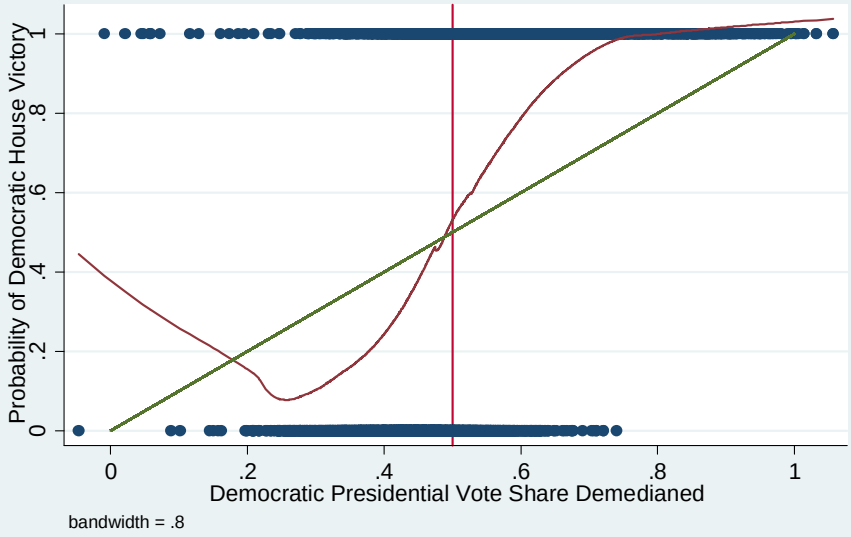

The lowess line in Figure 2 has the hypothesized relationship around 50 percent; the lowess line is much steeper than the green reference line. A one point increase in Democratic presidential vote share increases the probability that a Democrat will win the House seat in that district by more than one percent. Yet, the lowess line also predicts that a representative in a district with a demendianed democratic presidential vote share of 0

is more likely to go democratic than a district at .2. This is a spurious result caused by a few outlier cases in a section of the graph where there are only a few cases to begin with. Looking at a graph of a running mean smoother in Figure 3, the line is essentially flat from 0 to .2. Looking at lowess graphs by year in Figure 4, it is clear that the initial lowess graph is the spurious result of the blowout presidential election years 1964 and 1972 and the following midterms of 1966 and 1974 (which use the presidential vote share of the previous election) in which voters engaged in widespread ticket splitting. In 1964, this result appears only in the South, where LBJ did very poorly in a number of districts, but the Democratic house candidate still won the seat. In 1972, Nixon did unusually well in a number of districts across the country that nevertheless went Democratic at the congressional level.

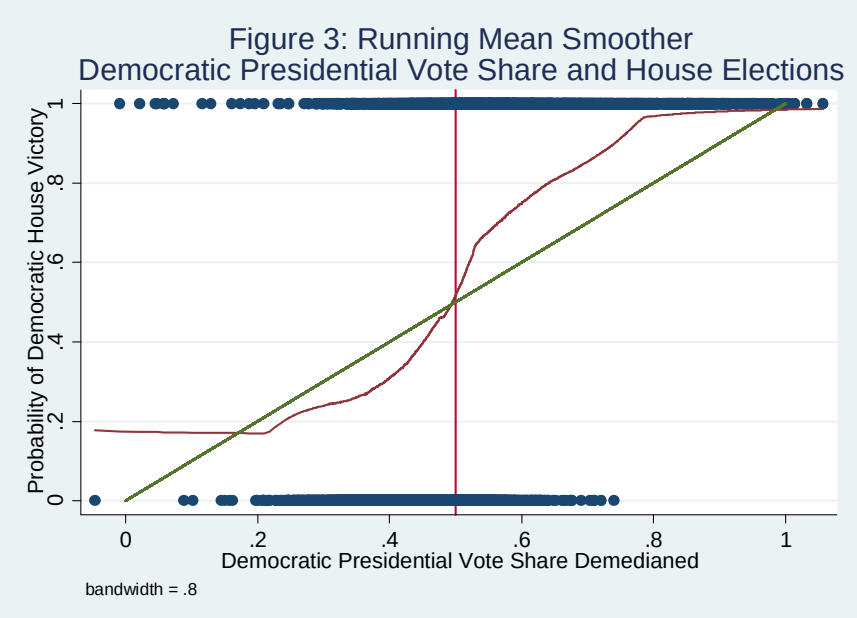

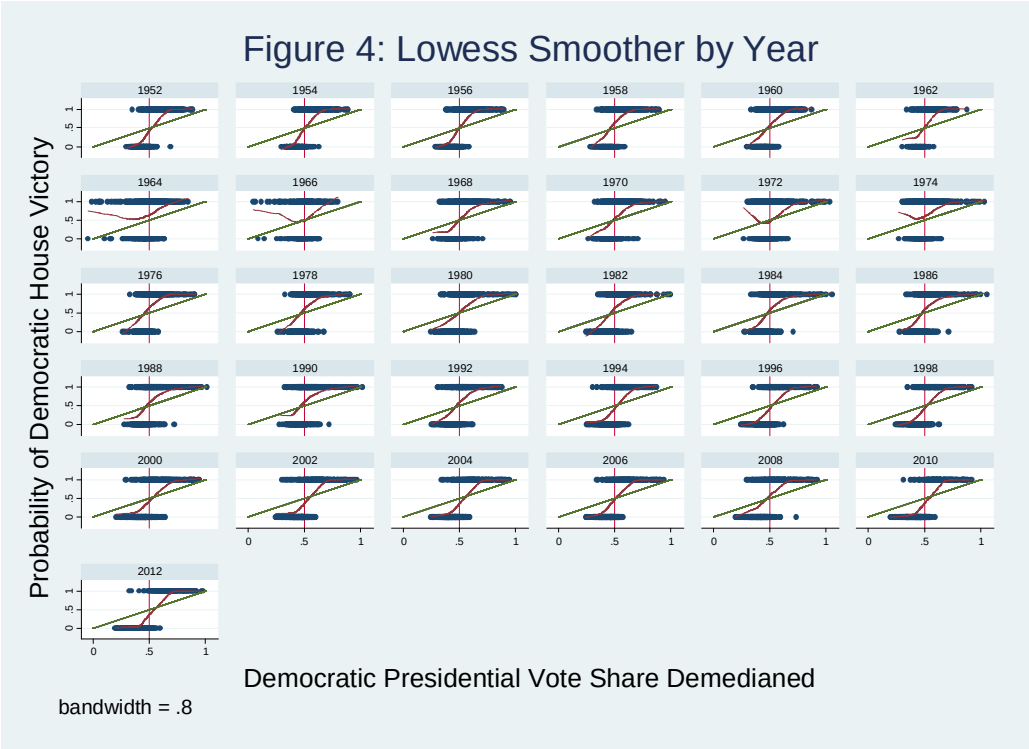

When blowout election years are excluded in Figure 5, the predicted s-shaped curve appears without spurious results at the tail ends. As predicted, a one point increase in democratic presidential vote share amounts to a more than one point increase in probability of democratic victory. A Democratic house candidate wins more than half the time in a 50/50 district because of incumbency.

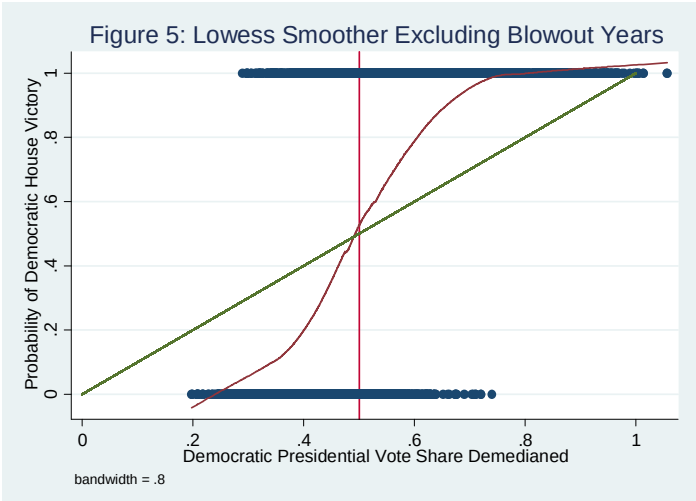

As Figure 6 shows, when races with incumbents are excluded, each party has an equal chance of winning if the demedianed Democratic presidential vote share is .5. Two party presidential share is thus a good proxy for district preferences. The shift between the point on the lowess line where the probability of a Democrat winning the seat is .5 in Figure 5 and where it is .5 in Figure 6 can be thought of as the incumbency effect.

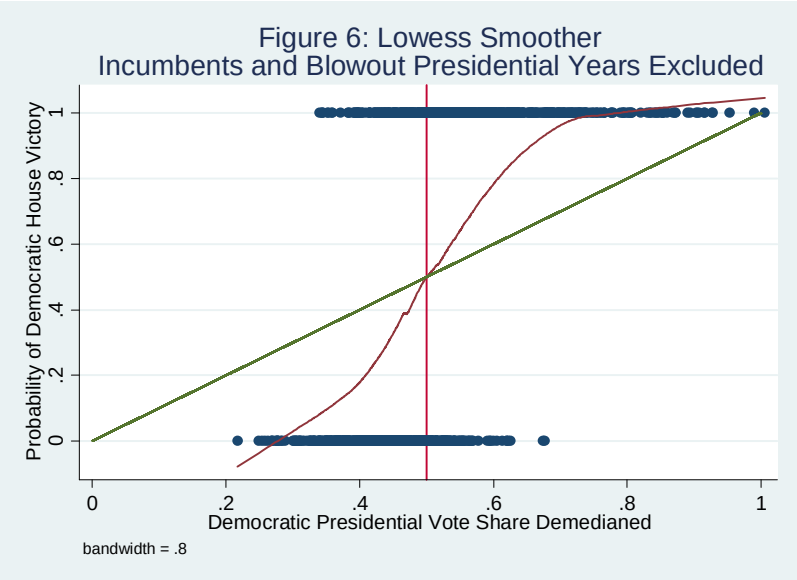

As the Democrats held the majority for most of the years examined, this point is left of Democratic presidential vote share equal to .5 (see red reference line). In years in which the Republicans held the majority, the lowess line reaches .5 probability after .5 Democratic presidential vote share demedianed, as can be seen in the lowess graphs by year in Figure 4.

The kernel density estimate of the Democratic presidential vote share and the lowess graph of the relationship between the presidential vote share and House election outcomes in a district demonstrate that Democratic presidential vote share is positively skewed across districts and that the electoral impact of this skew should in fact have a larger disparate impact on the Democratic Party than the nominal skew value calculated

by Mayhew. As Figure 6 shows when incumbency effects are removed, presidential votes share appears to be a good measure of district partisan preferences.

Table 2. Democratic Presidential Vote Share and Democratic Victory in a House District

Year	*Correlation*
1952	0.702
1954	0.701
1956	0.660
1958	0.566
1960	0.565
1962	0.585
1964	0.239
1966	0.167
1968	0.541
1970	0.518
1972	0.317
1974	0.296
1976	0.479
1978	0.469
1980	0.469
1982	0.516
1984	0.506
1986	0.494
1988	0.487
1990	0.460
1992	0.546
1994	0.627
1996	0.643
1998	0.646
2000	0.625
2002	0.688
2004	0.712
2006	0.685
2008	0.627
2010	0.758
2012	0.801

As you can see in Table 2, there is a strong correlation between the outcome of a House election and democratic presidential vote share. In the blowout elections of 1964 and 1972, as well as the subsequent midterms calculated based on the presidential vote share by district from the prior presidential election, the relationship between presidential vote share and house election results is relatively weak. As such, the results from these years

should be treated with additional skepticism. In the data set that I am working from, half of the cases are missing from both 1962 and 1966, so while I go ahead and report results for those years based on the cases that I do have, these year specific results are of little value.

The strength of the correlation between House election results and presidential vote share in a midterm election is not significantly different from the correlation in a general election. So while using presidential vote share from two years prior is not ideal, this indicates that it works just as well for predicting House election results two years hence as for the general election House results themselves. Finally, the relationship between the two variables is much stronger prior to 1964 and after 1994 than in the intervening years. This result is consistent with a story of southern partisan realignment starting in 1964 and ending in 1994.

4.1) Running Mean Smoother by Year

The advantage of using a running mean smoother, as opposed to a logit or probit, to determine the relationship between Democratic presidential vote share in a district and whether the Democratic house candidates wins that district is that it does not assume the functional form of the relationship between the two variables. This is important because a large portion of my argument hinges on this relationship being the predicted s-shaped form. So in order to apply a modified version of Gudgin and Taylor's (1980) method, which will be introduced in further detail below, I use a running mean smoother to non-parametrically map the relationship between presidential vote share in a district and which party won the House election in that district. Specifically, in order to assign a

probability of a House election victory conditional on presidential vote share in a district, for each year I create bins every tenth of one percent that take the average election outcome in the 5% bandwidth around that point for that year.[7] The smoother thus calculates the probability of victory in the actual election given a particular presidential vote share such that when a simulated district's presidential vote share is generated, it can be assigned the same probability of victory that the actual election results of that year predict.[8] No other independent variables are included in the running mean smoother, such as incumbency, because changing the distribution of presidential vote share across districts would also significantly alter incumbency. If a party has an incumbent advantage, it will be reflected in the results of the running mean smoother as incumbents will be able to win more often than non-incumbents otherwise would given the presidential vote share in a district. The running mean smoother thus does not seek to control for a party's incumbency advantage but allows it to be incorporated into the estimated relationship between district preferences and House election outcomes. While there is considerable normative concern about and empirical research on the impact of the incumbent advantage on the Democratic process, for the purposes of this paper it is simply a given that is indirectly reflected in the running mean smoother.

For midterm elections, the presidential vote share from the previous election is used to predict the outcome of each house seat. While not ideal because it assumes that

[7] This bandwidth was chosen because when districts with the Democratic presidential vote shares of the actual election districts are run through the election simulation model described in the following section, the actual election results are most accurately reproduced with a small bandwidth.

[8] This is true with one exception, House districts in which the Democratic presidential candidate did over 30 points worse than in the median House district, but the Democrat still won were dropped for the purposes of calculating the running mean smoother. This prevents the model from predicting that a district in which the Democratic presidential candidate did absolutely abysmally will still go to the Democratic House candidate. This only impacts a few cases in 1964 and 1966 and biases against a finding of electoral skew that favors the Republican Party because it underestimates the probability of an extremely conservative district electing a House Democrat.

district preferences do not shift vis-à-vis other districts in the intervening two years, having a separate running mean smoother for each year does allow the relationship between presidential vote share and house election results to vary between years. As the presidential vote share in a given district from a general election is just as strongly correlated with who wins in that district two years later in the midterm as it is with who wins the House seat during the same general election, presidential vote share is a workable proxy for preferences in midterm elections as well as general elections.

4.2) Modified Gudgin and Taylor Method

Gudgin and Taylor (1980) compare a party's actual seat share in parliament to what the results would have been with the same mean vote share averaged by district and the same standard deviation, but normally distributed. The difference between the two is the impact of the non-normality of the actual distribution. As I have discussed, missing cases make this method more difficult to apply in the American case and there is a theoretical reason to prefer to focus on the preferences that underlie House vote share, so I apply the method indirectly by considering what the results would have been with the same mean presidential vote share averaged by district and the same standard deviation, but normally distributed. Thus, I measure the impact of the non-normality of the distribution of preferences that underlies voting in House elections.

For each election year, I simulate 1,000 House elections in which the Democratic presidential vote share calculated by district is normally distributed across districts. For each simulated election, 435 House districts are drawn from a distribution with the same mean and standard deviation of Democratic presidential vote share as the actual House

districts, but which is normally distributed. The bins calculated with the running mean smoother of the actual House election results by year described in the previous subsection are then used to calculate the predicted probability for each simulated district that the Democratic house candidate will win given the Democratic presidential vote share in that district. Each district's predicted probability of Democratic victory is then compared to a random draw from a uniform distribution between 0 and 1 that simulates election conditions. If the Democratic probability of victory exceeds this draw, the Democrat won the district, if not the Republican is considered the victor. The results of all 435 districts can then be averaged to obtain the predicted seat share. This process is repeated 1,000 times, the mean of the simulations is reported as the predicted seat share under a normal distribution in Table 3. This seat share can be compared to the Democrat's actual seat share in the House to determine the impact of the non-normality of the actual distribution of preferences. Repeating the simulation 1,000 times allows us to provide bootstrap confidence intervals for the estimate.

The results of the simulations presented in Table 3 show that if preferences (as measured by Democratic presidential vote share) were normally distributed across districts, the Democratic Party would have done systematically better across almost all years. The simulation predicts that under a normal distribution of preferences across districts, the Democrats would have never lost control of the House of Representatives in the entire period from 1952 to 2008, and would have reclaimed the House in 2012. 1964 is the only exception in which the Democratic Party fairs better under the actual distribution of preferences as opposed to a normal distribution. Caution is in order in interpreting the results for blowout presidential election years like 1964 and 1972, as well

Table 3. Modified Gudgin and Taylor Measure of Electoral Skew

Year	*Actual Seat Share*	–	*Predicted Seat Share Under Normal Distribution*	=	*Electoral Skew*	*Predicted Majority Party Under Normal Distribution*
1952	49.08		54.29		-5.21 **	**1**
1954	53.10		57.79		-4.69 **	1
1956	53.79		57.94		-4.15 *	1
1958	64.03		66.43		-2.40	1
1960	60.18		62.13		-1.95	1
1962	56.62		59.89		-3.27 †	1
1964	67.82		67.56		0.26	1
1966	57.40		58.65		-1.25	1
1968	55.86		58.72		-2.86	1
1970	58.62		60.97		-2.35	1
1972	55.76		59.36		-3.60 †	1
1974	66.59		68.79		-2.20	1
1976	67.13		68.64		-1.52	1
1978	63.68		65.16		-1.48	1
1980	55.86		58.08		-2.22	1
1982	61.84		62.55		-0.71	1
1984	58.16		60.94		-2.78	1
1986	59.31		61.45		-2.13	1
1988	59.77		62.15		-2.38	1
1990	61.52		63.59		-2.07	1
1992	59.45		61.27		-1.82	1
1994	47.00		50.65		-3.64 †	**1**
1996	47.70		51.41		-3.72 †	**1**
1998	48.62		52.51		-3.89 †	**1**
2000	48.96		53.01		-4.05 *	**1**
2002	47.24		51.52		-4.28 †	**1**
2004	46.54		50.83		-4.29 *	**1**
2006	53.56		57.91		-4.35 *	1
2008	59.08		62.00		-2.92 †	1
2010	44.37		48.42		-4.05 *	0
2012	45.92		51.45		-5.53 **	**1**

Rusults reported in percentage terms of Democratic seat share on a scale of 0 to 100 (as opposed to 0 to 1) for facility of discussion. **Bold** entries indicate that predicted majority party is different than the actual majority party. *** = 99% confidence interval, ** = 95% confidence interval, * = 90% confidence interval, † = 80% Confidence Interval

as their subsequent midterms, because the relationship between Democratic presidential vote share and house election results is weak. Indeed, when the non-parametric bins used for these years are plotted the relationship between the presidential vote share of a district and the outcome of that district's House election does not inflect around 50 percent, but is

convex. Similar simulation results are obtained, however, when a logit is used in place of the non-parametric running mean smoother to give the relationship between presidential vote share and house election outcomes the assumed s-shaped functional form. The difference between the actual election results and the predicted election results under a normal distribution is smaller in 1964 than in any other year. In contrast, the difference in every year since 1994 in which the Republicans have held the House has been greater than in any year since 1956, averaging roughly an 18 seat bias in favor of the Republican Party.

The statistical significance of the electoral skew reported in Table 3 is based off of bootstrapped confidence intervals from the repeated simulation. As the confidence interval is two sided, a result that is significant at a 90% interval is interpreted as the actual election seat share being less than the simulated seat share in at least 95% of the 1,000 simulated elections. The predicted majority party under the normal distribution is reported in the next column. For each year in which the Republican Party in fact held a majority of seats, a seat share of 50 percent is within a 95% confidence interval of the predicted seat share under a normal distribution, so these results certainly do not indicate that in any given year between 1994-2004 or 2010-2012 the Democratic Party would have been guaranteed victory under a normal distribution of preferences. However, the consistent flip of predicted control suggests that the impact of the non-normality on the probability of House control is sufficient to have a substantively important impact on the probability of controlling the Speaker's gavel that has likely determined the majority party in at least some election years. This particular question will be addressed in more detail in the following subsection.

These results show that the impact of electoral skew is much larger than the nominal difference between the Democratic presidential vote share nationwide and the Democratic presidential vote share in the median congressional district that is reported by Mayhew. Indeed, the average reported impact is over 2.5 times larger than Mayhew's nominal measure. These results also fail to lend support to Rodden and Warshaw's (2009) supposition that that the skewed distribution of preferences makes it easier for the Republican Party to obtain a simple majority, but harder for them to obtain large majorities. It is easier for the Republican Party to obtain a simple majority, but to the extent that the model finds that the Democratic Party had an undue advantage in 1964, that advantage amounted to a single seat. As I have acknowledged, the low correlation between Democratic presidential vote share and House election outcomes in blowout election years makes 1964 difficult to model. However, even if a better measure of district preferences in this blowout presidential election were available and it showed a somewhat larger benefit to the Democratic Party in 1964, it would remain minimal in comparison to the advantage of the Republicans in the far more common close elections. The average advantage from the skewed distribution of district preferences accrued by the Republicans since 1994 is over 17 seats, a large enough advantage to flip control of the chamber in every instance except 2010, which is the only year in which the Republican Party is predicted to control the House under a normal distribution of preferences.

One skeptical response to these results could be to point out that the model predicts Democratic House victories even in years when the Republicans won a majority

of the House popular vote in the actual election; this seems too good to be true.[9] I would respond with an explanation of why this result is plausible given the model's assumptions by presenting a vignette of a simple spatial model of what occurs when preferences are redistributed normally across districts. Under the actual distribution of preferences, there are too many districts that are highly liberal for the distribution to be normal. Indeed, the distribution is skewed such that there are far more extremely liberal districts than extremely conservative districts. For a normal distribution, there must be more districts closer to 50 percent and fewer liberal districts out at the tail end of the distribution. Thus, some of the voters in highly liberal districts (and crucially far more voters than in highly conservative districts) must be redistricted into new districts, with much closer elections, for there to be a normal distribution of preferences across districts. So consider the voters of an actual congressional district in which the House seat was contested and a liberal Democrat (D_L) beat a Republican (R) with 70% of the vote. If the district was done away with and each of these voters was individually plopped into their own 50/50 House district where the president also received his national mean of 50%, who would they vote for? Well, that depends on the ideal points of the new candidates. The running mean smoother used in the statistical model calculates the probability of victory for a particular level of presidential vote share based on the actual candidates. So when a House election with a normal distribution is simulated, it is assumed that conditional on the presidential vote share drawn for a simulated district in a particular year, the candidates of the simulated districts are exactly like the actual candidates. Thus, the Democratic candidate in the new 50/50 district is considered to be a moderate Democrat (D_M) akin to the

[9] In comparison to my results, Gelman and King's (1991) results seem too good to be true. In any given year for the U.S. House in the Non-South from 1946-86, they find that there is a 5 to even 15 percent bias in favor of the Republican Party once incumbency effects are removed.

moderate Democrats who ran in 50/50 districts in the actual election. So assuming that the Democratic presidential vote share of the candidates have the following properties $D_L > D_M > R$, then all 70% of the voters who voted for D_L will also vote for D_M, but some of the remaining 30% who preferred R to D_L will now prefer D_M to R.[10]

Thus, the simulation of a normal distribution is not predicting that the Democratic Party would consistently win despite the Republican Party consistently winning the House popular vote, but that given the change to a normal distribution the House popular vote would also change, presumably enough to give the Democratic Party a popular vote majority as well as a majority of seats. How candidate ideology, voting behavior, and turnout would change given a different set of districts is obviously an empirical question, but it is beyond the scope of this paper and the statistical model is thus forced to make assumptions about these factors in order to capture the impact of non-normality. I present this vignette of the voters of a liberal district each being redistributed to their own brand new 50/50 district not to justify my assumptions, but to show why the results of the normal distribution simulation are perfectly plausible given my assumptions.

4.3) A Conservative Test for Change in Control of Congress

One possible criticism of these findings is that if I calculated the two party presidential vote share using the national popular vote rather than averaging presidential vote share across congressional districts, I would get a very different result. Gudgin and Taylor (1980: 519) recommend using average seat proportion, as I have done using the

[10] Notice that for simplicity sake R does not change in this little vignette. If in the actual district the Republican was a liberal Republican (R_L) and in the new hypothetical district they were a moderate Republican (R_M), then assuming that for Democratic presidential vote share $D_L > D_M > R_L > R_M$, the same story still holds true. Also note that while this story describes a single House district, it is meant to describe what the model assumes is going on at the aggregate level.

indirect form of average presidential vote share by district, to capture impact of "the non-normality of the actual distribution of constituency proportions." The primary focus of this paper is on how the skew in preferences across congressional districts affects election results, not on the impacts of turnout or malapportionment. For the purpose of measuring skew alone, presidential vote share averaged across districts is the correct input because it ignores turnout and malapportionment effects. Each district in a simulated House election is drawn from a distribution with the same mean and standard deviation for the presidential vote share as the actual congressional districts used to create the running mean smoother that gives each simulated district a predicted probability of a Democrat winning the House seat. Thus, when the probability of winning an actual congressional district is compared to the hypothetically normally distributed House, the difference in the results is entirely due to the non-normality of the actual distribution.

Therefore, if all this paper seeks to accomplish is to demonstrate the bias generated by the positive skew of democratic presidential vote share across congressional districts, it has done so. However, I also seek to make broader conclusions about how control of Congress would likely change under a fairer set of districts. As I acknowledge, in order to capture the effect of the non-normality of the actual distribution, I ignore turnout and malapportionment effects. If a different party is predicted to control the House under a normal distribution than the actual non-normal distribution, this provides a context for the extent of electoral skew. It does not, however, demonstrate that control of Congress would change if these districts were actually drawn up. By ignoring malapportionment and turnout, it assumes that all districts have the same size population and would have the same turnout rate under a different set of electoral circumstances.

While the necessary assumptions for measuring electoral skew due to the non-normality of Democratic presidential vote share across districts, using the national popular vote offers a set of more conservative assumptions that bias against a finding that control of the House likely changed hands under a different set of districts with more normally distributed preferences.

While using presidential vote share averaged across districts treats districts as if they all have the same population and everyone voted[11], using the national presidential popular vote assumes that a new set of districts would be apportioned by the number of people who actually voted, not the number of eligible voters or the population of the district. As Democrats are overly concentrated in urban districts with low turnout, this set of assumptions should favor the Republican Party (Rodden and Warshaw 2009). Indeed, the national Democratic presidential vote share calculated as the popular vote is ¾ of a percentage point less than the national Democratic presidential vote as averaged by district. Thus, for the purposes of predicting whether control of the House would switch from the Republican Party to the Democratic Party under a different set of districts with normally distributed preferences, using the national presidential popular vote share makes conservative assumptions that should bias us against a finding that control of the House would change hands from the Republicans to the Democrats.

Using the same modified Gudgin and Taylor method, I simulated 1,000 House elections for each year with that year's national two party vote share as their mean and the same standard deviation as previously used with the presidential vote share mean calculated across congressional districts. Because the District of Columbia is included in

[11] By everyone voted I mean that the results from the people who voted are imagined to be the results of everyone voting. If the model simulated universal turnout, the House Democrats would likely do slightly better (Citrin, Schickler, and Sides 2003).

the popular vote, but not the United States Congress, its vote totals were excluded from the calculation of the popular vote in order to have comparable populations. Given that the District of Columbia is overwhelmingly Democratic, this is also a conservative assumption.[12] Results are reported in Table 4.

Calculating the probability that a Democrat will win a seat using the same simulation model by year, the results from the national popular vote model are more mixed. From 1954 through 1960 and from 1976 through 1982, the model predicts the Republicans would do better under the popular vote mean normal distribution than the actual election results. However, from 1952 to 1982 the difference between the actual election results and the popular vote simulation is never enough to even come close to changing the outcome of an election for either party and on average the Republican Party still does better in the actual election. From 1984 onward, however, the Democratic Party always does worse in the actual election than in the popular vote simulation. In the critical period of interest from 1994 onwards, the Democrats pick up roughly 2.3 percent more seats in the simulation, or about 10 seats in the House. One potential reaction to this result is that 10 seats is a relatively small number of seats. This may well be an accurate read when one party controls a large majority in the House. However, in 1994, 1998, 2000, 2002, and 2012 the difference between the actual election results and the simulation based on the popular vote is more than the difference between a Republican majority and a Democratic majority.

In 1994 and 2002 the actual election result is outside of an 80% confidence interval for the simulated result and in 2012 it is outside of the 90% confidence interval.

[12] As we are considering how partisan electoral outcomes would change under a fairer distribution of preferences across districts, it would be altogether fitting to redistribute D.C.'s voters across congressional districts in which voters are not disenfranchised in House elections.

Like the results for a normal distribution based on presidential vote share averaged across districts, however, in the years where the

Table 4. A Conservative Test of Changing Party Control

Year	*Actual Seat Share*	−	*Predicted Seat Share Under Normal Distribution by Popular Vote*	=	*Difference in Means*	*Predicted Majority Party Under Normal Distribution by Popular Vote*
1952	49.08		49.45		-0.37	0
1954	53.10		52.89		0.21	1
1956	53.79		51.67		2.12	1
1958	64.03		61.35		2.68	1
1960	60.18		58.75		1.43	1
1962	56.62		60.33		-3.71 †	1
1964	67.82		68.01		-0.19	1
1966	57.40		58.02		-0.62	1
1968	55.86		56.20		-0.34	1
1970	58.62		58.85		-0.23	1
1972	55.76		59.68		-3.92 *	1
1974	66.59		69.32		-2.73 †	1
1976	67.13		65.59		1.53	1
1978	63.68		62.38		1.30	1
1980	55.86		54.84		1.03	1
1982	61.84		61.60		0.24	1
1984	58.16		59.75		-1.59	1
1986	59.31		60.28		-0.96	1
1988	59.77		60.72		-0.95	1
1990	61.52		62.20		-0.68	1
1992	59.45		60.97		-1.52	1
1994	47.00		50.35		-3.34 †	**1**
1996	47.70		49.20		-1.51	0
1998	48.62		50.33		-1.71	**1**
2000	48.96		50.26		-1.30	**1**
2002	47.24		50.24		-3.00 †	**1**
2004	46.54		48.99		-2.44	0
2006	53.56		55.90		-2.33	1
2008	59.08		60.33		-1.24	1
2010	44.37		46.40		-2.04	0
2012	45.92		50.38		-4.46 *	**1**

Rusults reported in percentage terms of Democratic seat share on a scale of 0 to 100 (as opposed to 0 to 1) for facility of discussion. **Bold** entries indicate that predicted majority party is different than the actual majority party. *** = 99% confidence interval, ** = 95% confidence interval, * = 90% confidence interval, † = 80% Confidence Interval

Republican Party in fact held a majority of the seats, a 50/50 election outcome is never outside of a 95% coincidence interval so I cannot definitively state that the Democratic Party would have won for sure in any given year. However, one advantage of running multiple simulations is that it can give us not just a prediction of which party would control the Speaker's gavel under a certain set of conditions, but also how likely that would be to occur. While all of the simulations run must necessarily be based on the assumption that the actual outcome was the most likely to occur in that particular election given the actual disticts, running multiple simulations can give us a sense of how close the winning party was to losing, even in the actual election. To do this, I reran the model not with draws from a normal distribution, but the actual 435 House districts in a given year.[13] Thus, using the Democratic presidential vote share in the actual districts, the same bins from the running mean smoother by year were used to assign each district a probability of a Democratic House victory in that district. The same procedure of drawing from a uniform distribution 0 to 1 was then used to determine the winner in each simulation of the actual districts. The percentage of the time that the Democratic Party was the majority party in the House as a whole under the simulation of the actual distribution, a normal distribution, and a normal distribution with the popular vote used as the mean is reported in Table 5.

Again, the simulation of the range of possibilities necessarily assumes that the actual result of the election was the mean of all possible results for that election given the actual districts. Given this assumption, however, even under the actual distribution there was a 1 in 5 chance of the Democratic Party retaking the House in 1998 and a 3 in 10

[13] This is also how I tested whether the model was unbiased. The Republican Party does a negligible amount better in the simulation of the actual districts than in the actual election, which biases the model against my hypotheses.

chance in 2000. Yet, in other years when the Republicans actually won control of the House, the odds of Democratic victory under the actual distribution were slim. In contrast, even under the conservative assumptions of using the popular vote model, the Democrats are much more likely to control the chamber. In any given year from 1994-2004 and 2010-2102,

Table 5. Probability of a Democratic Majority

Year	*Normal*	*Popular Vote Normal*	*Actual Election*
1952	96.7	40.6	29.1
1954	99.9	88.6	97.2
1956	100.0	75.8	99.0
1958	100.0	100.0	100.0
1960	100.0	100.0	100.0
1962	100.0	100.0	99.9
1964	100.0	100.0	100.0
1966	99.9	100.0	93.9
1968	100.0	99.7	99.9
1970	100.0	100.0	100.0
1972	100.0	100.0	99.6
1974	100.0	100.0	100.0
1976	100.0	100.0	100.0
1978	100.0	100.0	100.0
1980	100.0	97.4	100.0
1982	100.0	100.0	100.0
1984	100.0	100.0	100.0
1986	100.0	100.0	100.0
1988	100.0	100.0	100.0
1990	100.0	100.0	100.0
1992	100.0	100.0	100.0
1994	62.5	55.8	5.2
1996	71.7	36.2	11.0
1998	85.3	55.8	20.1
2000	91.2	54.1	29.4
2002	74.8	53.7	3.9
2004	62.9	33.4	1.4
2006	100.0	99.5	98.3
2008	100.0	100.0	100.0
2010	26.3	5.4	0.0
2012	71.9	56.0	0.0

the years in which the Republican Party won a majority seat share in the actual election, the Democratic Party averaged a less than 9% chance of winning control of the House in

the simulation based on the actual districts, but had a roughly 44% chance of winning under a popular vote based normal distribution, and a 68% chance of winning under a normal distribution with mean presidential vote share calculated by district average. So while we are only talking about an average difference of about 10 seats over this period, this difference would increase the Democratic Party's odds of controlling the U.S. House of Representatives by 35% in absolute percentage terms and almost 500% in relative terms. So while I cannot make a definitive claim about a change in the outcome of any one election, working under a set of conservative assumptions biased against my hypothesis, I find that under a different set of districts that were more fairly apportioned, the underlying probability of victory for the Democratic Party would have been substantially increased in close election years such that in at least some years the control of the U.S. House of Representatives has likely been determined by electoral skew.

5) Conclusion

The skewed distribution of preferences across districts is having a considerable impact on the composition of the U.S. House of Representatives. Even under a set of assumptions that were conservatively biased against the hypothesis, I find that the Republicans would have likely lost House elections in over half of the years between 1994 and 2004, as well as in 2012, under a normal distribution of preferences. The point of the paper, however, is not to make a definitive claim about any one election year, but to show that electoral skew is consistently having a negative impact on the House electoral prospects for the Democratic Party, particularly considering how close many recent House elections have been. Mayhew's (2011) measure of electoral skew, the

difference between the national presidential vote share and the presidential vote share in the median district, clearly established that there is an electoral skew that systematically favors the Republican Party. Yet, Mayhew downplays the impact of this skew on American politics. In this paper I show that when both the weight of the number of districts around the mean of the distribution of Democratic presidential vote share and the s-shaped relationship between it and a Democratic house candidate winning are taken into account, there appears to be a much more disparate partisan impact. On average, I find that the impact of electoral skew is 2.5 times larger than the impact that Mayhew reports. Mayhew (2011:13-14) argues that "the governmental system has bent towards *convergence* and *symmetry*," but where "significant impediments" prevented this from occurring, he "advance[s] the idea of *corrigibility*," or the idea that where there is a significant road block to policy progress, "given time, many such impediments are tackled and overcome." Based in large part on his finding that the skew in American legislative institutions is relatively small, he claims that "many alleged problems have proven to be nonexistent, short-term, limited, tolerable, or correctable" and as such "The public's shrug [to calls for reform] is understandable" (Mayhew 2011: 190). Given that neither major political party perceives itself to be cheated, the public's shrug is understandable. Yet, this paper demonstrates that the problem of electoral skew has both real and long term partisan consequences that consistently bias House elections against the Democratic Party. Mayhew (2011:179) points out that "occasional small statistical disadvantage has brought occasional brief grumbling, but little else has occurred. Yet a reform cause could crystallize quickly. As remarked in chapter 1, a triggering event

might be another 2000-type result that would arouse Democrats." However, 2000 was in fact another 1996-type event and it did not arouse a significant reform movement.

Little do most Americans know, 1996 was a year in which the Democratic Party won the popular vote in House elections, but remained in the minority in the House. In fact, in the wake of President Bill Clinton's reelection this fact seems to have gone entirely unreported. A search of the New York Times for the one year period following the election reveals no note in the paper of record that the Republicans retained a House majority despite more people voting to elect a Democratic congressperson than a Republican congressperson. Scholarly accounts share this glaring sin of omission, with Rivers, Cogan, and Brady (1997) titling an analysis of the election "The 1996 House Elections: Reaffirming the Conservative Trend." If neither journalists nor academics nor politicians tell people an undemocratic outcome has occurred, it is no wonder that the general public remains indifferent.

While the public's attitude is easily explained by elite behavior, elite indifference is more perplexing. To be sure, the pattern highlighted in the paper is not apparent with a simple glance at one year's election result. Nevertheless, the fact remains that there was a complete non-response from House Democrats in 1996, who had every reason to be just a little disgruntled. One possible explanation to explore is that given the Democrats had held a majority in the House for over 40 years and the problem had never arisen before, it seemed like a one-time event. In 2012, when the Democratic Party won the House popular vote but Republicans again retained control of the Speaker's gavel, there was some media coverage of the issue, but none of the coverage appears to have noted that this has happened before and is likely to happen again. Indeed, the media coverage in

2012 appears to have been short lived and focused primarily on Republican redistricting as the primary culprit.

While redistricting could certainly have played a part given that 2012 was the most partisanly skewed House election in the 1952-2012 period, the 2012 House election is part of a much longer trend that appears to mostly be a function of geography. Gudgin and Taylor (1974; 1976) argue that major left parties in general suffer from a positive skew in their electoral support under a system of single member plurality districts because their supporters are more tightly grouped together due to the political geography of cities. Rodden and Warshaw (2009), as well as Chen and Rodden (2013), offer strong evidence that the skewed distribution of preferences in the U.S. House of Representatives is in large part a function of geography. Whether districts are gerrymandered by politicians as in the American context or drawn up by an impartial committee as in the British context, it is extremely difficult to draw geographically contiguous districts that do not clump left wing supporters in one-sided urban districts. This appears to be a universal feature of single-member plurality districts: "Rydon (1957, p.61) has noted this phenomenon in Australia and New Zealand as well as in Britain and feels that it is 'an almost inevitable feature resulting from the concentration of Labour votes in industrial areas'" (Gudgin and Taylor 1974: 69). In the American case, electoral skew appears to have gotten progressively worse over the last 60 years as Americans, particularly liberal Americans, have become more and more likely to live in landslide counties (a geographic unit that is not gerrymandered) with like-minded partisans (Bishop and Cushing 2008). There have always been a disproportionate number of extremely liberal districts and that number was increasing long before 2012 (see Table 1).

Based on limited coverage in 2012 and the complete lack of news coverage in 1996, as well as the scholarly description of the 1996 election as part of a continuing conservative trend, it is unclear that even a blatant violation of democratic principles like the 1996 House election registers as an undemocratic outcome in the consciousness of most political elites, let alone ordinary Americans. The concept of a House popular vote, and the implicit understanding of what would happen under an alternative electoral system such as proportion representation, is certainly not a part of the American political lexicon. In contrast, the 2000 election provoked at least some discussion because there is a clear alternative to the Electoral College, i.e. the national popular vote. For most political elites and almost all Americans, electing a President by popular vote is perfectly conceivable. In contrast, electing a legislature by proportional representation based on a national popular vote, or even based on proportional representation in statewide districts, is a completely foreign idea to most Americans and does not appear on the radar screen of most political elites.

Even if the Democratic Party were more conscious of the problem, they would suffer a collective action problem in seeking reform because incumbent House members have an individual incentive to retain as liberal a district as possible. They would also have to overcome the considerable obstacle of explaining an alternative electoral system like proportional representation to the American people, let alone getting ¾ of the states to agree to amend the U.S. Constitution. All this makes a push for reform unlikely, particularly from House Democrats.

In order for a party to treat an election as an undemocratic outcome, it would seem that it must repeatedly be in the minority despite getting a majority of the popular

vote. In order for both parties to perceive a need for reform, the process must be blatantly arbitrary, i.e. in one year the Democratic Party wins the popular vote but loses the election and in the next year the Republican Party wins the popular vote but loses the election. Absent such blatantly biased or patently arbitrary outcomes, Mayhew is correct that a push for reform is unlikely to occur. And while what happened in 2012 has happened before and will likely happen again, it may well be over a decade or more before the Democratic Party again wins the House popular vote, but remains the minority party in the House. Whether or not Democrats remember what occurred in 1996, they seem condemned to repeat the events of 2012.

References

Brookes, R. H. (1960). The analysis of distorted representation in two-party single-member elections. *Political Science, 12,* 158.

Citrin, J., Schickler, E., & Sides, J. (2003). What if everyone voted? simulating the impact of increased turnout in senate elections. *American Journal of Political Science, 47*(1), pp. 75-90.

Dahl, R. (2003). *How democratic is the american constitution?*. New Haven: Yale University Press.

Hill, D., & McKee, S. C. (2005). The electoral college, mobilization, and turnout in the 2000 presidential election. *American Politics Research, 33*(5), 700-725.

Hopkins, D. (2009, September 3). *Whatever happened to moderate republicans? party asymmetry in the u.s. congress, 1972-2008.* Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1451202

Gelman, A., & King, G. (1994). A unified method of evaluating electoral systems and redistricting plans. *American Journal of Political Science, 38*(2), pp. 514-554.

Gelman, A., & King, G. (1990). Estimating the electoral consequences of legislative redistricting. *Journal of the American Statistical Association, 85*(410), pp. 274-282.

Gelman, A., & King, G. (1991). Systemic consequences of incumbency advantage in u.s. house elections. *American Journal of Political Science, 35*(1), pp. 110-38.

Grofman, B., Koetzle, W., & Brunell, T. (1997). An integrated perspective on the three potential sources of partisan bias: Malapportionment, turnout differences, and the geographic distribution of party vote shares. *Electoral Studies, 16*(4), 457-470.

Gudgin, G., & Taylor, P. J. (1980). The decomposition of electoral bias in a plurality election. *British Journal of Political Science, 10*(4), pp. 515-521.

Gudgin, G., & Taylor, P. J. (1974). Electoral bias and the distribution of party voters. *Transactions of the Institute of British Geographers,* (63), pp. 53-74.

Hill, D., & McKee, S. C. (2005). The electoral college, mobilization, and turnout in the 2000 presidential election. *American Politics Research, 33*(5), 700-725.

Johnston, R., Rossiter, D., & Pattie, C. (1999). Integrating and decomposing the sources of partisan bias: Brookes' method and the impact of redistricting in great britain. *Electoral Studies, 18*(3), 367-378.

Kendall, M. G., & Stuart, A. (1950). The law of the cubic proportion in election results. *The British Journal of Sociology, 1*(3), pp. 183-196.

Mayhew, D. (forthcoming 2011). *Partisan balance: Why Political Parties Don't Kill the U.S. Constitutional System* Princeton: Princeton University Press.

Rivers, D., Cogan, J. F., & Brady, D. (1997). *The 1996 house elections: Reaffirming the conservative trend* Hoover Institution. Retrieved from http://www.hoover.org/publications/monographs/27157

Rodden, J & Warshaw, C. (2009, February 24). *Whatever happened to moderate republicans? party asymmetry in the u.s. congress, 1972-2008.* Retrieved from http://politics.as.nyu.edu/docs/IO/10542/rodden_paper.pdf

Samuels, D., & Snyder, R. (2001). The value of a vote: Malapportionment in comparative perspective. *British Journal of Political Science, 31*(4), 651.

Taagepera, R. (1973). Seats and votes: A generalization of the cube law of elections. *Social Science Research, 2*(3), 257-275.

Tufte, E. R. (1973). The relationship between seats and votes in two-party systems. *The American Political Science Review, 67*(2), pp. 540-554.

Wildgen, J. K., & Engstrom, R. L. (1980). Spatial distribution of partisan support and the Seats/Votes relationship. *Legislative Studies Quarterly, 5*(3), pp. 423-435.