

The League of Women Voters of Pennsylvania et al v. The Commonwealth of Pennsylvania et al.

Nolan McCarty

Susan Dod Brown Professor of Politics and Public Affairs

Chair, Department of Politics

Princeton University

I. Introduction

I was retained by the legal firm Holtzman, Vogel, Josefiak, Torchinsky PLLC on behalf of their clients, the Speaker of the Pennsylvania House of Representatives, Michael C. Turzai and the President Pro Tempore of the Pennsylvania Senate, Joseph B. Scarnati, III, to provide expert opinions in connection with the matter *The League of Women Voters of Pennsylvania et al. v. The Commonwealth of Pennsylvania et al.* I am paid \$400 per hour plus costs for my services rendered in connection with this matter. My opinions expressed are in no way contingent on the payment of any monies owed to me for my services, and the monies owed to me are in no way contingent on the outcome of this matter. Specifically, my report responds to claims by Petitioners' experts relating to the purported extent to which the 2011 Pennsylvania congressional district map was designed to provide an undue advantage to Republican congressional candidates. My opinions expressed in this report are given within a reasonable degree of professional certainty.

II. Jowei Chen Report

In his report, Professor Jowei Chen provides an analysis based on 1,000 random simulations of Pennsylvania congressional districting plans. The first 500 of these simulations

are designed to produce districts that are relatively compact, contiguous, and have equal populations. He also suggests that his algorithm avoids county and municipal splits, but it is less clear how these criteria are achieved. I will refer to these results as Simulation 1. The second set of 500 simulations is designed to add considerations of incumbency by minimizing the extent to which the residences of incumbent members would be placed in new districts or that multiple incumbents might have to run for the same seat.¹ I will refer to these results as Simulation 2. I will not comment directly on the methodology behind Professor Chen's simulations. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

Given the size of these sets, we cannot be confident that the 500 districting plans in each of Professor Chen's simulation are at all representative. This problem is exacerbated when Professor Chen looks at the even smaller number of simulations that generate at least one majority African-American district.

My main focus will instead be on the way that Professor Chen measures partisanship of the enacted districting plan and that of his simulated plans. As discussed on pages 12 and 13 of his report, Professor Chen uses the results of the Pennsylvania statewide elections from 2008 to 2010.³ He then classifies districts in the enacted and simulated plans as Republican or Democrat depending on which party received the most votes across the six elections. His primary

¹ Since Pennsylvania lost a seat through reapportionment, completely avoiding multi-incumbent elections would have been impossible.

³ Those election include Presidential, Attorney General, Auditor General, and State Treasurer elections in 2008 and the US Senator and Gubernatorial elections in 2010.

empirical claim is that the number of Republican districts in the enacted plan (13) is larger than the number of Republican districts in almost all of his simulated plans. He concludes from this that the enacted plan was drawn to favor the Republican Party.

The primary problem with his approach to measuring partisan bias is that whether a district cast a majority of its ballots for Republicans in statewide elections is a very imperfect indicator of how the district will vote in actual congressional elections. For example, Professor Chen's approach suggests that nine of Pennsylvania's nineteen districts in its 2004 plan were Republican. Yet over the course of the decade that plan was in place, the number of Republicans in the Pennsylvania House delegation ranged from 7 to 12 with an average of 10.2. So not only does the actual number of Republican-controlled districts vary markedly from Professor Chen's prediction over time, but the prediction was wrong on average. My key point is that if his measure is a poor predictor of Republican seat shares, then the performance of both the enacted plans and the simulated plans will be distorted. Thus, it is impossible to conclude that the expected results of the enacted plan diverge from those of the simulated plans.

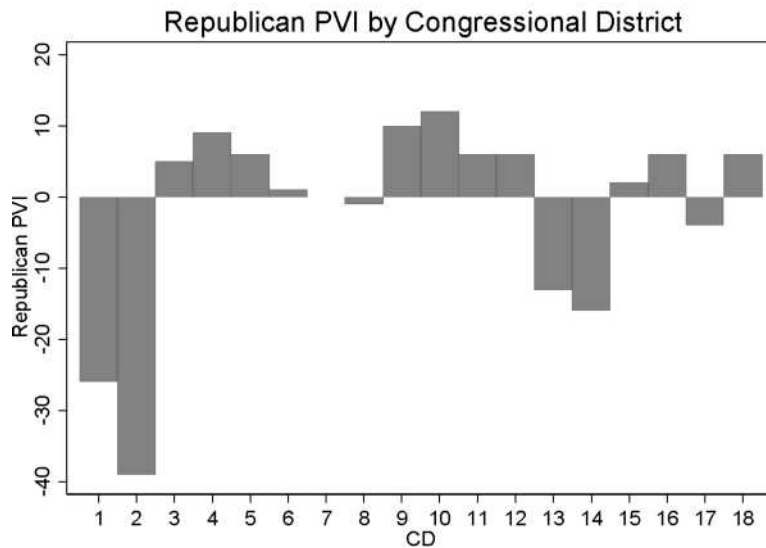
To demonstrate the magnitude of the limitations of Professor Chen's measure, I will use a similar measure of district partisanship that is available for all US House districts over time. That measure is the Partisan Voting Index (PVI) developed by Charles Cook.⁴ For congressional districts, the PVI is simply the average partisan presidential vote in the district over the past two elections minus the average partisan presidential vote nationally over the same period. The national average is subtracted to adjust for any candidate effects and national swings. A PVI of D+2 means that a district's voters were two percentage points more favorable to the Democratic presidential candidate than the nation as a whole. A PVI R+3 means the district voted three

⁴ <http://www.cookpolitical.com/pvi-0>.

percentage points more for the Republican than the country. For presentational purposes, I rescale the PVI so that Republican leaning districts are assigned positive numbers and the Democratic leaning districts are assigned negative numbers.⁵ When scaled in this way, I refer to the measure as the Republican PVI.

First, I compute the Republican PVI of the 2011 congressional map at the time of its adoption using the 2004 and 2008 presidential elections. Based on those results that I report in Figure 1, there are 6 Democratic-favorable districts, one even district, and 11 Republican-favorable districts. Note that using a slightly different voting measure to capture district partisanship, I estimate only 11 Republican districts compared to Professor Chen’s estimate of 13. But importantly, there is no a priori reason to prefer one measure over another. Since they correlate at .999, each must be an equally good (or bad) predictor of actual House elections.

Figure 1



⁵ In other words, a R+3 is assigned 3 and a D+2 is assigned -2.

But even Figure 1 may overstate the extent to which the 2011 enacted map favors the Republicans. While a Republican may be more likely to win election in a R+1 district than a Democrat is, a Republican victory is by no means assured. In fact, even under normal electoral circumstances, Democrats are likely to win a significant number of such Republican-leaning districts. So any reasonable prediction of the partisan performance of a districting plan should incorporate the fact that either party has a good chance of winning Republican-leaning competitive districts such as Pennsylvania districts 6, 7, 8 and 15 as shown Figure 1. In fact, as I show below, historically Democrats have been able to win a reasonable proportion of congressional districts that are as Republican as districts 3, 4, 5, 16 or 18.

To demonstrate, I use data on presidential and congressional elections to relate the Republican PVI of a district with the likelihood that it was won by a Democrat. In so doing, I use data for all congressional elections from 2004 to 2014. The congressional election data comes from Gary Jacobson, a leading political scientist, whose data is widely used in academic work on congressional elections.⁶ Then to compute the PVI measure, I use data from PoliData, an electoral data firm, that I obtained in 2014 for an unrelated academic project.⁷ For each congressional election, the Republican PVI is computed as described above: the average Republican presidential vote share in the district over the previous two elections minus the national Republican presidential vote share. I round these calculations to the nearest whole number for ease in presentation.

With measures of PVI and congressional election outcomes in hand, my analysis is simple. For each value of PVI, I compute the proportion of times a Democrat won in such a

⁶ See “JOPrepfile1” <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/29559>.

⁷ McCarty, Nolan “Reducing Polarization: Some Facts for Reformers” *The University of Chicago Legal Forum* 2015: 243-278.

congressional district over the 2004 to 2014 period. For example, I find that Democratic candidates won 39.6% of the 63 elections held in R+1 congressional districts. Moreover, they won only 46.4% of the 44 elections held in D+1 (i.e. Republican PVI = -1) congressional districts. We can think of these proportions as a rough estimate of the probability that a Democrat can win a congressional seat with a particular partisan composition.

Figure 2 summarizes my findings for all values of Republican PVI between -9 and 9.⁸ The circles indicate the relative number of elections used to estimate each proportion. The line is a loess line that shows a locally-weighted prediction of the actual win share as a function of the PVI. First, note that the probability of a Democratic win changes smoothly as the Republican PVI decreases. There is no dramatic jump at Republican PVI = 0. In fact, actual Democratic win rates are very similar at Republican PVI = 1 and Republican PVI = -1. Second, note that increasing the Republican PVI is associated with substantial changes in the Democratic win rate in districts as extreme as those with Republican PVI of -8 or 8.

⁸ A full table of all of these results including the number of elections in each category is including in the appendix to this report.

Figure 2

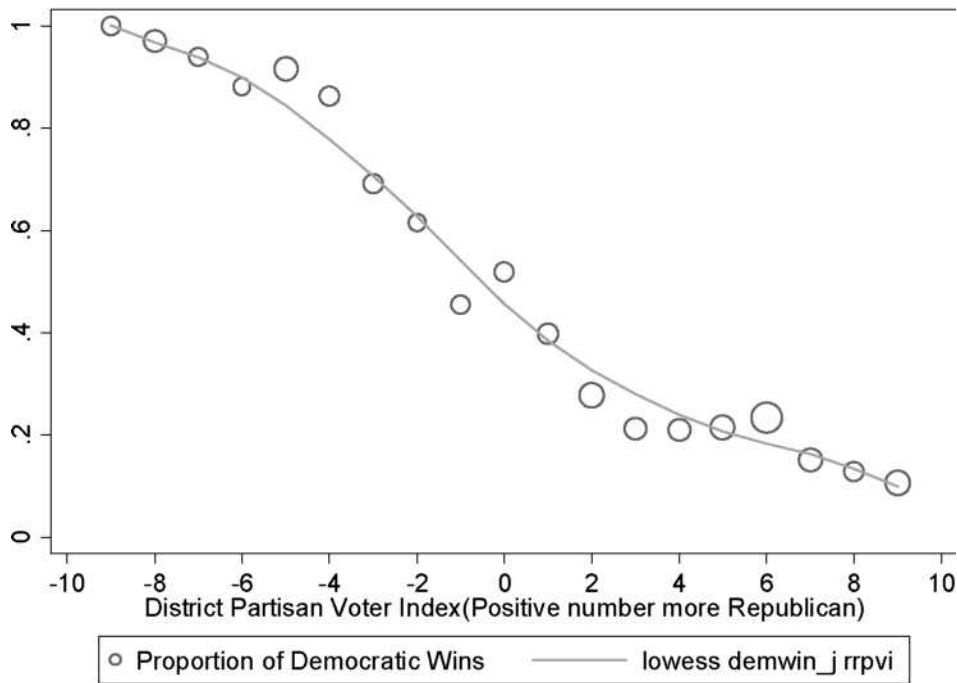


Figure 2 also reveals that several of the congressional districts labelled by Dr. Chen as Republican seats are ones that the Democrats should have a reasonable probability of winning. For example, even the R+6 districts have been won by Democrats about 23% of the time over the previous decade.⁹ Moreover, given the results displayed in Figure 2, it is clearly misleading to assess the partisanship of districts on a dichotomy based on whether the Republican PVI is greater than or equal to zero. A much better assessment would use the results of Figure 2 to calculate the proportion of congressional districts that each party should be expected to win. For comparison purposes, I estimate partisan performance of both the 2004 districting plan and the 2011 districting plan by assigning to each of the districts the probability that a Democratic candidate would win based on its PVI. Then by summing these probabilities, we can obtain the

⁹ It is important to note that this estimate is based on a very large number of elections. There were 129 elections held in R+6 congressional districts, more than in any other type of district. Eight districts in the enacted plan have Republican PVIs of 6 or less (districts 3,5,6,7,11,12,16, and 18).

expected Democratic seat share for each plan. Those results are reported in Table 1. For each plan I have computed Republican PVIs for each district. These PVI measures are based on the 2004 and 2008 presidential elections so that the measures are as comparable across the two plans as possible.¹⁰

Table 1 shows the results for each districting plan side by side. The left panel reports my estimate of the probability that a Democrat should have won each of the 19 districts in the 2004 plan given their partisanship as measured by PVI. If we sum those probabilities, we obtain the expected number of Democratic seats. The result is that the Democratic candidates should have won between 9 or 10 of the 19 seats, or obtain a seat share of .503. The right panel shows the results for the 2011 districting plan. Here I compute that the Democratic candidates should have won about 8 of 18 seats, or 45%. So while the 2011 plan is estimated to result in more Republican seats than the 2004 plan, the difference is only marginal. Based on my calculations, the number of expected Democratic seats fell by about 1.4 (from 9.55 to 8.15). If the 2011 map performed similarly to the old map in partisan terms, Democratic candidates would have been expected to win about 9 seats so a decline in the expected number Democratic seats of .55 is due to the state's loss of a congressional district following the 2010 Census. A loss of less than one seat, roughly .85, can be associated with the map.

¹⁰ The PVI for the current districts could be computed using the 2016 and 2012 election results. But for purposes of evaluating the changes to districting, it is much more relevant to have a measure of the PVI at the time that the new districts were drawn.

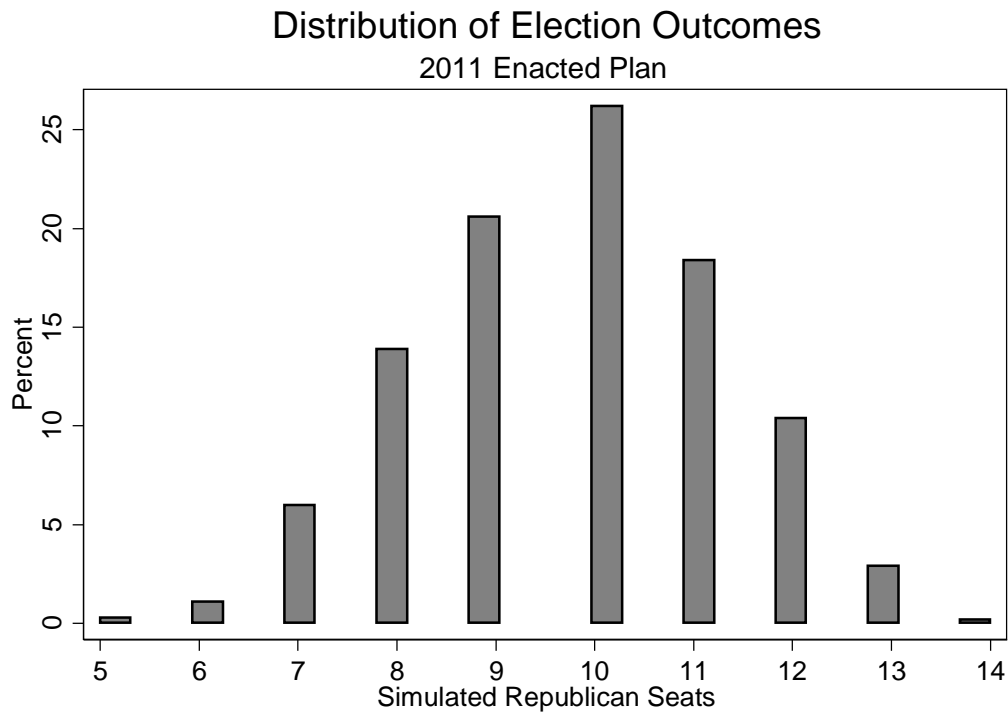
Table 1: PVI and Democratic Win Probabilities by CD						
2004 Congressional Districting Plan				2011 Congressional Districting Plan		
CD	PVI	prob(Dem)		CD	PVI	prob(Dem)
1	-35	1.000		1	-26	1.000
2	-38	1.000		2	-39	1.000
3	3	0.212		3	5	0.214
4	6	0.233		4	9	0.106
5	9	0.106		5	6	0.233
6	-4	0.863		6	1	0.397
7	-4	0.863		7	0	0.519
8	-2	0.615		8	-1	0.455
9	17	0.023		9	10	0.050
10	8	0.127		10	12	0.107
11	-4	0.863		11	6	0.233
12	1	0.397		12	6	0.233
13	-7	0.939		13	-13	1.000
14	-19	1.000		14	-16	1.000
15	-2	0.615		15	2	0.277
16	8	0.127		16	6	0.233
17	6	0.233		17	-4	0.863
18	6	0.233		18	6	0.233
19	12	0.107				
Expected Dem seats		9.555		Expected Dem seats		8.150
Expected Dem Share		0.503		Expected Dem Share		0.453

It is important to note that Table 1 indicates only the expected number of Republican seats under the enacted plan, but not how the number might vary given the uncertainty in election outcomes. To illustrate this variation, I use the probability of each district electing a Democrat for the 2011 enacted plan from Table 1 to simulate 1000 elections.¹¹ For each of these simulations, I can compute how many seats were won by the Republicans, those results are presented in Figure 3. Note that across the 1000 simulations, the number of Republicans elected varies from as few as

¹¹ For example, from 2004-2014 Republican candidates won R+2 districts 72.3% of the time. So in the simulations, R+2 districts are won by Republicans with probability .723.

five to as many as 14, but the expected number is around 10.¹² Note that the outcomes where the Republicans win 13 seats in the enacted plan are relatively rare for the *given partisan configuration of the 2011 enactment*. That the Republicans have won 13 seats under the enacted plan suggest that they have over performed or that Democrats have underperformed relative to historical standards. Had the outcome of 13 seats been due to an intentional partisan gerrymander, one would expect that it would be a more frequent outcome in my simulated elections. Thus, my results suggest the 13 seats currently held by Republicans cannot be attributed to the partisan features of the map.

Figure 3



¹² Because I have assumed that each election outcome is statistically independent, Figure 3 probably underestimates the variation in election outcomes. If I had assumed that outcomes within a single simulation were positively correlated, one would expect more extreme outcomes.

This analysis has two important implications for evaluating Professor Chen's report. First, his measure dramatically overstates how favorable the 2011 enacted plan was to Republicans. Based on historical patterns, such a plan should have produced from 9 to 11 Republican seats. Second, his simulations also ignore the historical relationship between district partisanship and congressional election outcomes. Similar to Figure 3, we should expect each simulation to produce a variety of outcomes, including extreme ones with more than 13 Republican seats, due to the uncertainties associated with election outcomes. Thus, his partisan performance of simulations may well be consistent with that which we observe under the enacted plan.

These measurement problems extend to Professor Chen's analysis of the output of his simulations. Like his measurement of the enacted plan, his measurement of the partisanship of simulated plans ignores electoral uncertainty – a district simulated at R+1 is called a Republican district even though such districts are often won by Democrats. To illustrate the magnitude of these problems, I use Professor Chen's measure of partisanship (Republican vote share from 2008-2010) for each district in each of his simulations. I then use regression analyses to convert these measures into Republican PVI so that I can apply my estimates of the Democratic win rate (e.g. Figure 2 and the appendix).¹³

Just as I did above for the 2011 plan, I can compute expected number of Republican seats for each of Professor Chen's simulations. Across the 500 simulations in his first set, I compute that the average expected number of Republican seats was 10.97 with a standard deviation of .229. Note that when I incorporate election uncertainty, the expected number of Republican

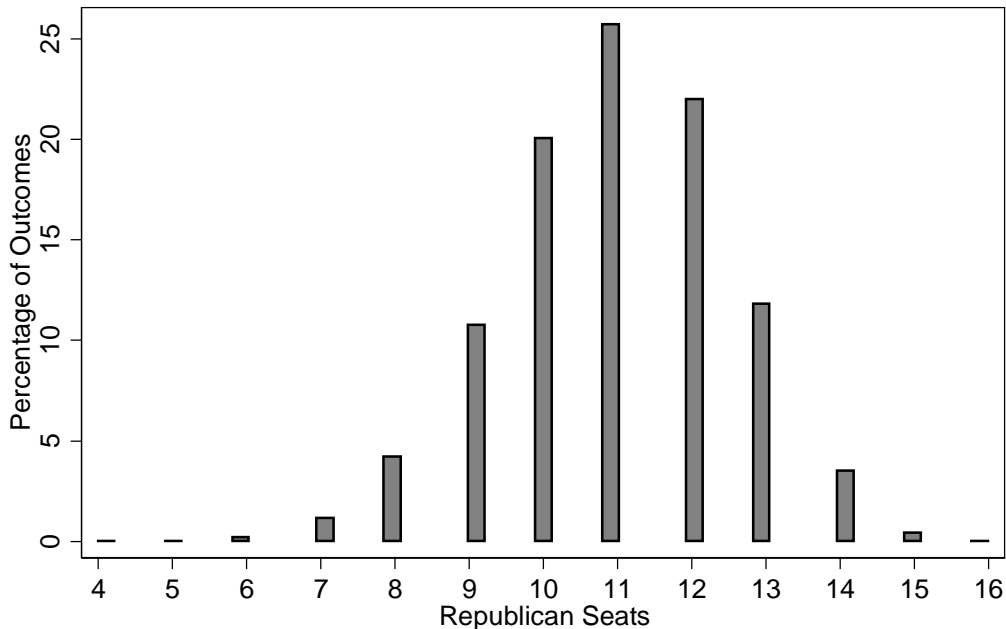
¹³ Chen's Republican vote shares were converted into Republican PVIs by regressing the PVI on the Chen vote share for the 19 congressional districts under the 2004 plan. This regression suggested that $PVI = 96.532 * (\text{Chen vote share}) - 44.926$. The R^2 of this regression was .998.

seats is considerably higher than the typical 9 reported by Professor Chen. But more importantly it is larger than the 9.85 expected seats in the enacted plan. In fact, the least Republican simulation had 10.3 Republican seats. Once the uncertainty about election outcomes is incorporated, ALL of Professor Chen’s simulations in set 1 are more favorable to Republicans than the 2011 enacted plan.¹⁴

Additionally, I can examine how the actual number of Republicans seats across Professor Chen’s simulations by simulating election outcomes under each as I did above in Figure 3. Now for each of the 500 simulations in Simulation 1, I simulate 500 election outcomes using the probability from the table in the appendix. This exercise produces 2500 different observations of Republican seat shares. The distribution of these outcomes is show in Figure 4.

Figure 4

Outcomes of Simulated Elections in Simulated Plans
Chen Simulation 1



¹⁴ Analysis of simulation 2 reveals similar patterns. In that simulation, the average expected number Republican seats was 11.2 with a standard deviation of .267. The minimum expected number was 10.54.

One take away from Figure 4 is that a very wide range of outcomes can be observed from the simulated districts. There are simulations where the Republicans win as few as 4 and as many as 16 seats. Second, if we compare Figure 4 to Figure 3 we see more evidence that the simulated plans are more favorable to Republicans than the enacted plan. Thirteen Republican seats was a relatively rare outcome in the enacted plan, but more than ten percent of the simulations produce such an outcome.

In summary, once the uncertain mapping between district partisanship and congressional election outcomes is taken into account, Professor Chen's simulations provide no evidence that the 2011 enacted congressional map is an "outlier" with respect to its partisan advantages.

III. Gerrymandering and Polarization

In his report, Christopher Warshaw argues that the impacts of gerrymandering on congressional representation are especially pernicious now that the two political parties in Congress have polarized. As the parties have polarized, Republicans and Democrats have tended to pursue increasingly divergent policy interests. Because of this partisan divergence, voters who do not support the party of the district's representative are likely to be poorly representative.

Polarization does indeed pose challenges for good representation of voter interests and for effective legislative governance. Yet Professor Warshaw's efforts to link the effects of polarization to gerrymandering are not compelling. First, while he is careful not to suggest otherwise, there is very little evidence to suggest that gerrymandering is an important cause of polarization. The US Senate is increasingly polarized despite fixed state boundaries. Increasing polarization is evident in the delegations of small states where districting is less important, and

indices of polarization do not rise disproportionately following redistricting cycles. Statistical evidence and simulations also fail to find much of any effect of gerrymandering on legislative polarization.¹⁵

Therefore, given the lack of evidence on the causal link between gerrymandering and polarization, Professor Warshaw focuses on the claim that gerrymandering exacerbates the problems associated with polarization. The key to his argument is that Republican-biased gerrymanders leave a disproportionate number of Democratic voters in districts with Republican representatives whose voting in Congress are not likely to reflect the interests and preferences of Democratic voters. He writes that

The growing pro-Republican Efficiency Gap creates conditions where many Democratic voters in Pennsylvania are unable to elect representatives of their choice. The growing polarization in Congress means that representatives in Congress nearly always vote the party line. So Democrats that are artificially deprived of the opportunity to elect someone that shares their values [and] do not have their views represented in Congress. This means that they have little, if any, voice on important issues. Thus, the combination of partisan gerrymandering and polarization in Congress has a profound, pernicious effect on democratic representation. (p. 15)

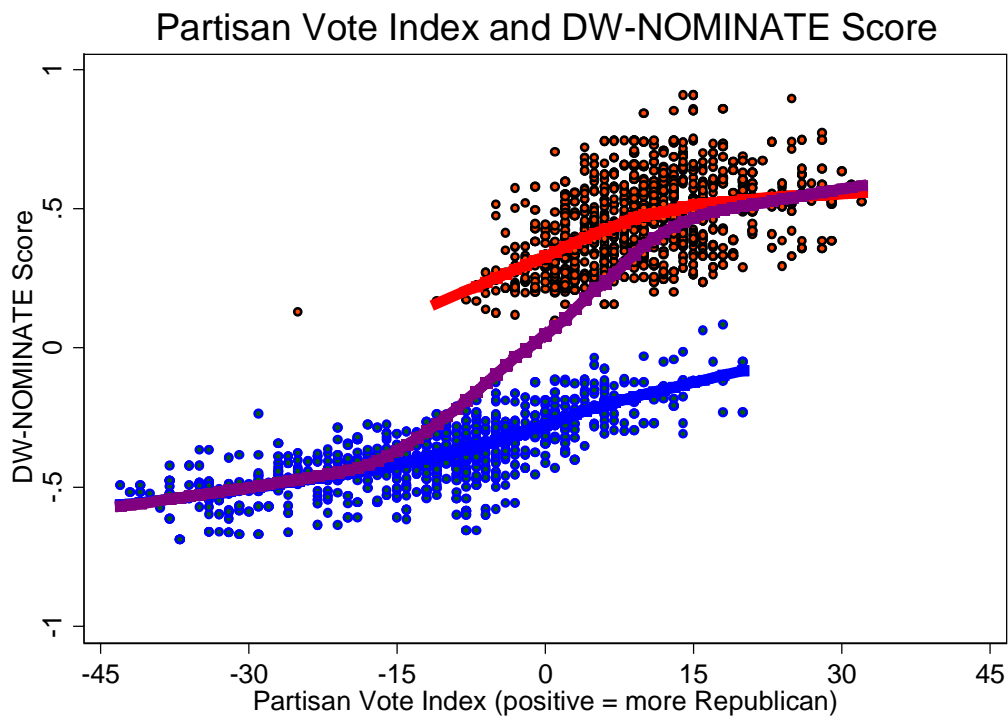
There are two reasons to be skeptical of this argument. First, under any alleged Republican gerrymander, Democrats will not just be cracked but also packed. Thus, Democrats who are placed in heavily Democratic districts will have an extraordinary opportunity “to elect someone that shares their values.”¹⁶ Moreover, Democrats elected from districts with large Democratic majority may be much more likely to vote for the policies favored by Democratic voters as they would have to appeal to very few Republican constituents.

¹⁵ See McCarty, Nolan, Keith T. Poole, and Howard Rosenthal. 2009. “Does Gerrymandering Cause Polarization?” *American Journal of Political Science* 53(3):666-680.

¹⁶ Quoting Warshaw report, page 15

Second, under the logic of gerrymandered districts, “cracked” Democratic voters are likely to be placed in districts with relatively small Republican majorities. But as I demonstrated in the last section, Democratic candidates win Republican-leaning districts with some regularity. So these Democratic voters have some chance of being represented by members of either party. So in expectation, representation in these competitive districts will not be as extreme as those in safe districts where either Democratic or Republican voters are packed.

Figure 5



To quantify evidence for my concerns, I combine the data on district Republican PVIs with the DW-Nominate measure of the conservatism of each House member’s voting record – the measure used by Professor Warshaw. Figure 5 shows a simple plot of the PVI of each U.S. House district from 2004 to 2014 against the DW-Nominate conservatism score. The red dots

represent the positions of Republican House members while the blue dots represent those of Democrats.

First, note that the voting records of members of both parties get considerably more conservative as the Republican PVI increases. The solid lines for both parties represent the locally best fitting predictor of the DW-Nominate score given the district's PVI. For both parties, but especially the Democrats, these lines slope steadily upward. Thus, it is clear that the most liberal voting records are compiled by members representing the most Democratic districts. Moreover this relationship holds no matter how Democratic the district is. Thus, these results support the idea that Democratic voters, to the extent to which they would like their representatives to vote in a liberal way, may prefer to be packed into Democratic majority districts.

Second, consider those districts with Republican PVI's just above zero. Under Professor Warshaw's theory, these Democratic voters in marginally Republican districts should be poorly represented. Again it is not clear why this need be the case. The results of the previous section are apparent in Figure 5 – a lot of Democrats have represented Republican-leaning districts over the past decade. So at least some of the time, Democratic voters in these districts do get represented by one of their co-partisans. A second effect is also apparent. Republicans who represent competitive districts are more moderate than those in safe Republican districts. Thus, Democratic voters in Republican-leaning districts do appear to have some influence over the positions taken by their Republican representatives. To aggregate these two effects, a locally-weight prediction line is included for the total relationship between PVI and the DW-Nominate score. The purple line shows the best prediction of the DW-Nominate score for a given PVI. Note that for Republican PVI's between -15 and 15, the *expected* voting record of a representative

is very responsive to the partisan composition of the district. Consistent with the points raised in the last section, there is not much difference in the expected position of a representative in a -1 district from that of a +1.

Polarization may indeed lead to poor representation if members of Congress generally have positions more extreme than those of voters. But the evidence in Figure 5 lends little support to the idea that the packing and cracking associated with gerrymander exacerbates the disconnection between voters and their representative.

IV. Efficiency Gap Analysis

Much of Professor Warshaw's report centers on an analysis of the so-called "efficiency gap." This measure is based on the notion of "wasted votes." The number of wasted Democratic votes is measured as the number of Democratic votes in seats won by Democrats beyond the 50% needed to win. In seats won by Republicans, all Democratic voters are considered wasted. Republican wasted votes are defined analogously.¹⁷ The efficiency gap is simply the absolute difference in the number of wasted Democratic votes minus the number of wasted Republican votes divided by the total number of votes. The intuition behind this measure is that if the excess wasted votes of the disadvantaged party were reassigned to other districts that party could have obtained greater legislative representation. While intuitive as a rough measure of partisan advantage, it is considerably less reliable in proving the existence of partisan gerrymandering.

¹⁷ The computation of the efficiency gap faces a significant measurement problem of how to deal with uncontested legislative elections. Thus, analysts are forced to make ad hoc assumptions such the minority party would have received 25% or to use statistical modelling as Professor Warshaw has done.

The primary problem with the analyses of the efficiency gap is that it is well-established that the efficiency gap can vary across states for reasons having very little to do with partisan gerrymandering. In particular, the geographic sorting of partisans can lead to large increases in the efficiency gap. Proponents of the measure have not developed principled ways of determining when an efficiency gap is too large to be justified by geographic sorting. Therefore, they rely exclusively on empirical standards noting when efficiency gaps are large relative to other states and to previous plans within a given state.¹⁸ But because the ingredients of partisan sorting vary across states and over time, such standards are not adequate to establish that a particular efficiency gap measure shows evidence of partisan gerrymandering.

Evidence provided in Professors Warshaw's report underscores these concerns. Consider Figure 2 which plots the national efficiency gap and Figure 4 which plots the historical efficiency gap in Pennsylvania. Clearly, the measure moves as much in between redistricting efforts as it does following the implantation of a new plan. Given this natural variation, it is difficult to draw any strong conclusions about the causal impact of redistricting from efficiency gap measures.

Another source of variation in the efficiency gap are how well parties perform in close elections. Suppose a Republican candidate wins a district with 51% of the vote. Then, 49% of the votes in that district are wasted Democratic votes while only 1% of Republican votes are wasted. But suppose the Democrat had won the seat by that same small margin. Then the rates of wasted votes would be dramatically reversed. As I documented in Section II, both of these outcomes are extremely likely for any district that is not too partisan. Consider a district with a

¹⁸ For example, see Warshaw Figure 5 and the surrounding discussion.

Republican PVI of +2. About 70% of the time, elections in that district will produce an excess of wasted Democratic votes but 30% of the time more Republican votes will be wasted. Thus, the efficiency gap can swing dramatically based on which way such a district lands. This problem is especially pertinent in terms of measuring the efficiency gap in Pennsylvania. As I noted in Section II based on national patterns of the last decade, Pennsylvania Democrats should have been expected to win about 45% of the 18 seats (roughly 8 seats) under the 2011 plan. Instead, they control five. Had they performed up to expectations under the enacted plan, the patterns of wasted votes in 3 seats would be dramatically different and the overall efficiency gap in Pennsylvania would be much lower. Importantly, this effect is not merely hypothetical. Note from Figure 4 of Professor Warshaw's report that under the 2001 plan, Pennsylvania's efficiency gap ranged from a 12% pro-Republican gap to a 4% pro-Democratic gap.¹⁹ The reason the gap shifted dramatically in the Democratic directions in 2004 and 2006 was because the Democrats were very successful in winning competitive seats.²⁰ Democratic candidates won in districts with Republican PVIs of 1, 3, 6 (twice), and 8. Several of these seats were lost in 2010 and the Democrats have fared poorly in comparable competitive Republican-leaning districts since then. Therefore, it seems very hard to attribute that entire lack of success to gerrymandering.

Finally, it is not clear conceptually whether surplus votes for the winning party and votes for the losing party should be considered "wasted." As I demonstrated with Figure 5 in Section III, the voting records of representatives are highly responsive to the partisan composition of the district. With the possible exception of the most extreme Republican districts, the voting records

¹⁹ Recall from Table 1 that these two plans had very similar partisan features.

²⁰ Despite this evidence of the instability of the efficiency gap, Professor Warshaw includes an analysis of the durability of the gap (section 4.5, page 11). In this analysis, he simply correlates the efficiency gaps in 2016 and 2012 by state. But a positive correlation of these measures within state is to be expected since features such as the geographic concentration of voters change very slowly over time.

Professor Warshaw argues that due to high levels of partisan polarization, partisan gerrymandering has large negative impacts on the representation of Pennsylvania Democrats. In response I show that his analysis ignores three factors: the representational benefits to Democrats from living in a strongly Democratic district, the effects of Democratic voters on the behavior of Republican representatives in competitive districts, and the opportunities for Democrats to win competitive congressional districts. These three factors taken together clearly reduce the negative effects on Democratic voters. Finally, I question Professor Warshaw's use of the efficiency gap as evidence for a Republican gerrymander. As I illustrate, this measure is subject to large amounts of variation across states and over time that cannot be attributed to partisan bias in districting. In particular, I argue that the efficiency gap in Pennsylvania is inflated by the historical underperformance of Pennsylvania Democrats in congressional elections.

Appendix

Republican PVI and Democratic Win Rates 2004-2014		
Republican PVI	Number of elections	Proportion of Democratic Victories
32	2	0.000
31	2	0.000
30	2	0.000
29	8	0.000
28	12	0.000
27	3	0.000
26	20	0.000
25	12	0.000
24	10	0.000
23	12	0.000
22	2	0.000
21	25	0.000
20	27	0.148
19	26	0.000
18	37	0.081
17	43	0.023
16	59	0.034
15	57	0.053
14	83	0.096
13	82	0.073
12	56	0.107
11	56	0.054
10	80	0.050
9	85	0.106
8	55	0.127
7	73	0.151
6	129	0.233

5	84	0.214
4	67	0.209
3	66	0.212
2	83	0.277
1	63	0.397
0	54	0.519
-1	44	0.455
-2	39	0.615
-3	55	0.691
-4	51	0.863
-5	72	0.917
-6	42	0.881
-7	49	0.939
-8	68	0.971
-9	44	1.000
-10	29	1.000
-11	48	1.000
-12	51	1.000
-13	53	1.000
-14	32	1.000
-15	40	1.000
-16	28	1.000
-17	17	1.000
-18	26	1.000
-19	26	1.000
-20	21	1.000
-21	27	1.000
-22	21	1.000
-23	27	1.000
-24	14	1.000
-25	14	0.929

-26	18	1.000
-27	21	1.000
-28	8	1.000
-29	17	1.000
-30	4	1.000
-31	14	1.000
-32	21	1.000
-33	8	1.000
-34	25	1.000
-35	20	1.000
-36	8	1.000
-37	6	1.000
-38	13	1.000
-39	1	1.000
-40	1	1.000
-41	6	1.000
-42	3	1.000
-43	3	1.000



December 5, 2017