Income, Ideology and Representation*

Chris Tausanovitch†
Department of Political Science
UCLA

September 2014

Abstract: Do legislators represent the rich better than they represent the poor? Recent work provides mixed support for this proposition. I replicate the result of differential representation using a dataset on the political preferences of 362,487 individuals, but I show that the extent of this differential representation is limited. In recent years, representation occurs primarily through the selection of a legislator from the appropriate party. Although the preferences of higher income constituents account for more of the variation in legislator voting behavior, higher income constituents also account for much more of the variation in district preferences. In light of the low level of overall responsiveness, differential responsiveness appears small.

*Prepared for the Russell Sage Foundation Journal of the Social Sciences, Conference on Big Data in Political Economy

†Assistant Professor, Department of Political Science, UCLA, ctausanovitch@ucla.edu

Few issues in the study of representation have garnered more attention in recent years than the link between economic inequality and political inequality. The majority view in this literature argues that government responds to the preferences of large segments of the population that are higher income much more then it responds to large segments of the population that are lower income (Bartels, 2009; Bonica et al., 2013; Butler, 2014; Gilens, 2012; Gilens and Page, 2014). This is finding implies that even though citizens are equal in their ability to vote, they are unequal in their ability to incentivize politicians to take particular policy positions. This difference remains even after accounting for different rates of voting (Bartels, 2009), begging the question of why legislators would ignore a substantial portion of the people to whom they owe their jobs and continued reelection.

The approach pioneered by Bartels (2009) is useful for answering this question. Bartels examines the dyadic relationship between legislators and different classes of constituents within their districts. He supposes that legislators vote according to the wishes of their mean voter, as per various "mean voter theorems" (e.g. Caplin and Nalebuff, 1991; Schofield, 2007). If this is the case, then each citizen should receive equal weight in a legislator's voting decision, just as a each quantity is weighted equally in the calculation of an arithmetic mean. Groups should be weighted in correspondence with the size of each group. Given a set of non-overlapping groups, this leads to a clear specification for determining whether legislators do weight groups equally: assume that legislator positions are a function of the positions of the mean of each group multiplied by it's relative size, but allow the actual weights to vary.

However, representation of the mean voter is not the only way that legislators may make their decision. The mean voter replaces the older and more prominent theory of representation of the median voter (Black, 1948; Downs, 1957). If legislators represent the

median voter, they may have equal regard for all of their constituents, but nonetheless some will seem to recieve more "weight" than others: in particular, those that vary more, or are otherwise more likely to determine the location of the median. Yet another way in which legislators may represent their districts is via the mediating effects of party (Campbell et al., 1966). Voters may simply choose a candidate of the party they prefer, and conditional on choosing the right party, legislators may be bound only by the standards of loyal behavior within their respective party. If legislators have this degree of latitude, then whichever group is most likely to determine the balance of party support within a district will appear to have the greatest weight, despite an equal regard for all voters by the legislator.

In this paper, I replicate the existing result of differential representation using a large dataset of political preferences. However I also show, following two recent papers using other data sources, that it is not robust (Erikson and Bhatti, 2011; Brunner, Ross, and Washington, 2013). The conclusion depends on whether the slope or the fit of the model is thought to be a better indicator of responsiveness. More importantly, I show that these models do not necessarily imply large substantive differences in legislator positions resulting from differential representation. In univariate models, all groups are substantially predictive of the position that legislators take.

In the current climate, legislator positions are relatively homogenous within parties and heterogenous across parties. This is in direct contrast to the mean voter theorem- the theoretical assumption underlying Bartels's (2009) empirical specification. Partisan theories of representation are a reasonable alternative that match better with this empirical reality. Perhaps groups are unequal in the extent to which they determine the party of their representative. This leads to a different empirical specification where party, not legislator position,

is the dependent variable (Brunner, Ross, and Washington, 2013). In the latter part of the paper, I show that this specification leads to a similar conclusion. Both the positions of low income and high income constituents explain the party of their representative reasonably well. The slope of this relationship is steeper for higher income constituents. Discriminating between different explanations for legislator behavior is a difficult task that is beyond the scope of this paper. Nonetheless, the fact that these two very different approaches give similar results is a useful starting point.

These results should be understood in the context of broader political realities. The correspondence between political preferences of constituencies writ large and the preferences of their representatives is remarkably weak, particularly when party is accounted for (Clinton, 2006; Tausanovitch and Warshaw, 2013). During a period when Congress has polarized dramatically, the distribution of preferences of the public has remained centrist and stable, highlighting this disconnect (Fiorina and Abrams, 2012; Hill and Tausanovitch, 2014). In other words, responsiveness at an aggregate level appears to be poor. This is not surprising in light of recent evidence that casts doubt on the notion that a significant number of voters choose candidates on the basis of policy (Tausanovitch and Warshaw, 2014). Furthermore, few voters are aware of the policy stances of their particular representatives (Tausanovitch and Warshaw, 2014).

In light of these findings, one might wonder whether disaggregating constituents into high and low income groups could provide an explanation. After all, if legislators only respond to high income constituents then "averaging in" lower income constituents will create the appearance of weaker representation. The results below show that this does not appear to be the case. Separating groups by income and introducing income itself as an additional

variable does not appear to substantially improve our ability to predict the positions that legislators take.

In the section that follows I explain the methodology I use to estimate policy positions. In the next section I describe the data underlying the analysis. Section 4 presents the results concerning mean voter representation, and section 5 presents the results regarding partisan representation, followed by the conclusion.

Measuring Preferences

One of the core difficulties in measuring policy preferences is that statements of preferences on individual issues may not accurately reflect underlying attitudes. Respondents may make top-of-the-head judgements based on immediately available considerations (Zaller, 1992) or their choice may be affected by purely idiosyncratic or irrelevant factors (e.g. Achen and Bartels, 2012). A solution to this problem is to aggregate preferences in some way. Multiple (putatively) independent indicators of political preferences are less affected by random noise than a single response. Research has shown that in the case of voters, using multiple indicators increases the predictive power of attitudes on outcomes such as vote choice (Ansolabehere, Rodden, and Snyder, 2008).

The most commonly used methods for measuring underlying positions from revealed preference data are item response models. Item response models conceptualize preferences as a continuous latent variable in an underlying preference space. Individual choices depend on the choosers' latent preferences and the features of that particular choice. One of the simplest cases is a one-dimensional quadratic utility binary response model (Clinton, Jackman, and

Rivers, 2004). Let x_i denote person i's latent ideology, and y_{ij} denote person i's response to question j, where $y_{ij} = 1$ indicates a "yes" response and $y_{ij} = 0$ indicates a "no" response to question j. Then, the probability that person i responds "yes" to question j is

$$Pr(y_{ij} = 1) = \Phi(\beta_j x_i - \alpha_j)$$

where Φ is the standard normal cumulative distribution function, and α_j and β_j are the item parameters for question j. β_j captures the direction of the item (is yes a liberal or conservative response) as well as how strong the relationship is between responses to the item and underlying preferences. α_j captures the underlying liberalism or conservatism of the item (how liberal does one typically have to be to respond yes/no). The model is identified by restricting the x_is to have mean 0 and standard deviation 1, and the direction is fixed so that negative values are liberal.

This simple model allows us to estimate preferences and take account of the fact that some questions are more informative than others in different parts of the preference space. I estimate this model using a Bayesian approach, with dispersed normal priors for each of the estimated parameters. Unfortunately it is quite computationally expensive to run in standard implementation. Using software developed with Jeffrey Lewis (UCLA), I parallelize a Markov Chain Monte Carlo estimate of this model using data augmentation. In each iteration of the Markov Chain, posterior draws from the distribution of the item parameters and person parameters can be drawn indepentently. We conduct these draws simultaneously on Graphical Processing Units, allowing us to achieve speeds 32 times faster than standard implementations of this model. Using this software, I am able to estimate latent preferences

for a dataset of 362,487 survey respondents, containing 5,084,676 non-missing responses to 264 items.

There are numerous advantages to using a continuous measure of political preferences based on responses to policy questions, but the most important one is that a continuous measure of preferences simply gives us more information about the location of individuals in the policy space. This may be the reason that Erikson and Bhatti (2011) are able to find differential representation using the 9,253 respondents to the NES Senate study with a 7-point measure of ideological self-placement, but are unable to find differential representation using the 155,00 respondents to the NAES with a 5-point measure. The less granular measure does not distinguish as well. Ideological self-placement is particularly poorly suited to studying differences between representation among income groups because income groups are not likely to have a shared understanding of the question. This is in addition the benefits of using multiple items to mitigate measurement error.

Data

Analyzing representation requires data on the policy preferences of constituents and the legislators who represent them. Data on the former come from 9 large-sample political surveys: the 2000 and 2004 National Annenberg Election Survey (NAES), and the 2006, 2007, 2008, 2009, 2010, 2011 and 2012 Cooperative Congressional Election Studies (CCES). Combined, these studies provide responses from 362,487 individuals over this period of 12 years, for an average of 833 respondents in each of the 435 congressional districts.

The household income of respondents is self-reported on each survey in a series of cate-

gories. In the CCES, there are 16 income categories. However, the categories differ across the NAES and CCES. When categories are consolidated into 6 groups, they perfectly coincide. To classify the income of groups within congressional districts, I further consolidate income into four groups: those making less than $25,000, those making $25,000-$49,999, those making $50,000-$99,999, and those making more than $100,000. I call these groups "low income," "medium-low," "medium-high," and "high income" respectively. 78% of respondents chose to answer the income question, reducing the sample to 282,701. 20% of respondents are classified as low income, 27.5% medium-low, 33.6% medium-high, and 16.5% are classified as high income.

There are 264 unique policy questions in this dataset. However, responses to these questions are sparse due to the fact that different surveys ask different policy questions. Following Tausanovitch and Warshaw (2013), I identify the positions of respondents to different surveys relative to one another by constraining common questions to have the same item parameters. In addition, I use smaller sample surveys attached to the large 2010 and 2011 CCES surveys to provide more linking questions. The purpose of these surveys was to ask 177 of the questions that had been asked in other surveys in order to estimate the item parameters in a common space. This method for linking large sample surveys, and the data for doing so, comes from Tausanovitch and Warshaw (2013).

Data for the positions of legislators comes from Poole and Rosenthal's DW-NOMINATE scores (Poole and Rosenthal, 1997)[1]. Although the functional form of DW-NOMINATE is different than the Bayesian quadratic item response model outlined above, in practice it results in very similar estimates, and so for convenience I use it here. Rather than respond to

[1]data provided at www.voteview.com, accessed in August of 2014

survey questions, members of Congress cast roll call votes on policy issues. DW-NOMINATE scores are calculated using members' roll call votes as their statements of preference. Since these votes are actually yes-no choices they are amenable to a binary model.

Districts are matched to their respective members of the House of Representatives for the 111th session of Congress (2009-2011) [2]. One possible source of discrepancy between this paper and prior work is the focus on the House, unlike prior work that focuses primarily on the Senate (Bartels, 2009; Erikson and Bhatti, 2011). The CCES provides district identifiers for each respondent. For the NAES, I match respondents probabilistically to their districts using their zip codes. In most cases, zip codes are fully contained within districts, but in cases where there is partial overlap with multiple districts, the extent of the overlap is used to calculate the probability that a given respondent resides in a given district. Districts are from the 2000 redistricting in order to match up with the 111th session.

Results: Income and Representation

Is it in fact the case that higher income voters are better represented than lower income voters? Although more attention has been given to articles arguing in the affirmative, there are some nicely executed counterexamples. Using much larger sample sizes than both the original Bartels (2009) study, both Erikson and Bhatti (2011) and Brunner, Ross, and Washington (2013) find mixed evidence of differential representation. Both studies have disadvantages. Erikson and Bhatti (2011) use respondent's self placement using abstract ideological

[2]The results are quite similar if we match districts to legislators from the 112th Congress. It is possible to conduct the analysis using multiple congresses at once, but this counts districts more than once. Choosing a single congress is a more conservative approach.

labels (very liberal, liberal, moderate, conservative, and very conservative) as the measure of respondent positions. Brunner, Ross, and Washington (2013) use ballot propositions to measure voter ideology. There are a limited number of such propositions in each election, and income is measured at the neighborhood level, in California only. Nonetheless, their data comes from the universe of voters and as a result their sample size is enviable. In contrast to these papers, I use a large national sample of individuals responding to large numbers of policy questions, with income measured at the individual level.

Table 1: Pearson correlations between mean preferences of incomes groups within congressional districts

	μ_H	μ_{MH}	μ_{ML}	μ_L
μ_{MH}	0.87			
μ_{ML}	0.84	0.90		
μ_L	0.75	0.81	0.80	
μ	0.92	0.97	0.95	0.86

N=435, in all cases p $<$ .001

Why have existing studies come to different conclusions regarding representation? One explanation is that the variables that we seek to distinguish are highly collinear, and measured with error. Table 1 shows the pairwise correlations between five variables. μ_H is the mean preferences of high income constituents (income $>$ \$100,000), μ_{MH} is the mean preferences of medium-high constituents (\$50,000-\$99,999), μ_{ML} is the mean preferences of medium-low constituents (\$25,000-\$49,999), and μ_L is the mean preferences of low income constituents ($<$ \$25,000). μ is the the mean preferences of all constituents. The lowest correlations are between the preferences of low income constituents and other groups, but all of these correlations are very high. Unlike past studies, these measures reduce error through the use of a measurement model. Nonetheless, all of the quantities are measured with error, in part due to measurement and in part due to sample size. High correlations between

each quantity raises the possibility of autocorrelated error, which can cause instability in regression coefficients.

I begin by replicating the methodology used by Bartels (2009). In each district I calculate the percent of the sample that falls into each group. I call this p_g where g indexes the group: L, ML, MH, and H, respectively. I then decompose the mean preferences in each group, multiplied by the proportion in that group. By Bartels' logic, if legislators represent mean preferences in their district without regard to income, then the coefficients in a regression of legislator position on the proportion-weighted group means should all be equal. If the coefficient on one group is higher than the others, this suggests a counterfactual where legislators change their positions more in response to this group than to the others.

Table 2 shows the results of three regressions. The first two are univariate models that simply regress the position of the legislator on the preferences of low-income and high income people, respectively. The third is the specification from Bartels (2009) that includes each group and weights them by their proportion in the district. By Bartels' criteria, this model refutes the hypothesis that the rich are better represented than the poor. On the contrary, if anything low income people appear to be better represented. Not only is the coefficient in model 1 for low income preferences greater than the coefficient in model 2, for the high income, but in the combined specification the lowest income group has the greatest coefficient. The effects in model 3 have an oddly non-linear pattern, with the poor garnering the greatest coefficient but the medium-low group recieving a coefficient that is indistinguishable from 0. This is contrary to any expectation from the literature, and to my own expectation. What could explain these results?

One possibility is that the proportion of rich and poor constituents are variables that

Table 2: Regression of legislator position on income group preferences

	Dependent variable:		
	Legislator DW-NOMINATE score		
	(1)	(2)	(3)
constant	0.27^{***}	-0.09^{***}	0.07^{**}
	(0.025)	(0.016)	(0.027)
μ_L	1.24^{***}		
	(0.074)		
μ_H		0.71^{***}	
		(0.039)	
$p_H \times \mu_H$			1.47^{***}
			(0.365)
$p_{MH} \times \mu_{MH}$			0.92^{***}
			(0.321)
$p_{ML} \times \mu_{ML}$			0.32
			(0.405)
$p_L \times \mu_L$			2.63^{***}
			(0.521)
Observations	435	435	435
R^2	0.393	0.425	0.523
Adjusted R^2	0.392	0.424	0.523
Note:		*p<0.1; **p<0.05; ***p<0.01	

capture the urban/rural split that we observe dividing districts represented by Democrats from those represented by Republicans. The variance of preferences among the poor is much lower than the variance of preferences among the rich, likely because of greater measurement error in the preferences of the poor. However, the proportion of the poor that cross a threshold of "liberalness" may be a good indicator of an urban district, a poor district, or a majority minority district. Rather than gather detailed district-level data, a simple way to account for this sort of possibility is to control for the proportions of the district sample that are in each income group. Table 3 does just this, replicating each column from Table 2 but with controls for the proportion high, low, and medium-low income, with medium-high as the excluded category. This specification is similar to the one used by Erikson and Bhatti (2011).

The results from Table 3 are much more intitutive than the results from Table 2, and are closer to previous findings. Controlling for the income of a district, legislator responsiveness appears to increase with the income of the each group. Nonetheless, these coefficients are not significantly different from one another.

Given the ambiguity of the results in Tables 2 and 3, it is too soon to conclude that the poor are dramatically underrepresented. Research tends more often than not to find that the poor are underrepresented, but better data has not increased our confidence in this conclusion. However, the coefficients in these regressions are not the only way one could conceptualize responsiveness or representation.

The approach taken to examining responsiveness so far assumes that the coefficient on group-level preferences is the best measure of whether legislators are "responding" to voter preferences. It is difficult to know how to interpret this coefficient. Legislator positions

Table 3: Regressions controlling for income-only variables

	Dependent variable:		
	Legislator DW-NOMINATE score		
	(1)	(2)	(3)
Constant	0.31	−0.19	−0.34
	(0.342)	(0.334)	(0.315)
μ_L	1.37***		
	(0.078)		
μ_H		0.76***	
		(0.041)	
$p_H \times \mu_H$			1.55***
			(0.355)
$p_{MH} \times \mu_{MH}$			1.33***
			(0.310)
$p_{ML} \times \mu_{ML}$			0.97***
			(0.400)
$p_L \times \mu_L$			0.84
			(0.560)
p_H	0.58	0.71	1.44**
	(0.603)	(0.592)	(0.559)
p_L	−1.77***	−1.51***	−0.69
	(0.485)	(0.477)	(0.456)
p_{ML}	0.91	1.05	0.91
	(0.784)	(0.767)	(0.714)
Observations	435	435	435
R^2	0.463	0.484	0.577
Adjusted R^2	0.458	0.479	0.570
Note:			*p<0.1; **p<0.05; ***p<0.01

and voter positions are measured quite differently, and so a larger coefficient could measure overreactions to constituent preferences as easily as it measures better representation. One simple question we might ask is whether the positions of the poor or the rich are more accurate predictors of legislative positions. If legislators are truly focusing on one group more than the other, then our predictions of legislator positions should be closer to the truth when we use the preferences of the more well-represented group as a regressor. The evidence on this question from Tables 2 and 3 is clear: variance explained is always higher using the high income group than the low income group.

To understand what this means exactly, consider Figure 1. This figure graphs the univariate regression line of legislator positions on the positions of high income constituents, overlaid on the scatterplot of the data. The grey lines in the figure show the regression lines for Republican legislators only (the top cloud) and Democratic legislators only (the bottom cloud). The reason for showing these regression lines should be clear from the plot. The relationship between legislator and constituent positions is hardly linear. The polarization in legislator positions means that the transition between liberal and conservative legislators is not smooth. In contrast, the positions of high income constituents are spread relatively smoothly throughout the preference space. The bottom line here is that most of the variance, and hence most of the variance explained, is between-party. Within-party the lines are relatively flat and the variance explained is much less.

Figure 2 shows the univariate regression line, scatterplot, and associated within-party regression lines when the positions of low income constituents within districts is used as the explanatory variable. There are two main differences between this plot and the previous one. First of all, the variance explained is lower, both within and between party, while the slope

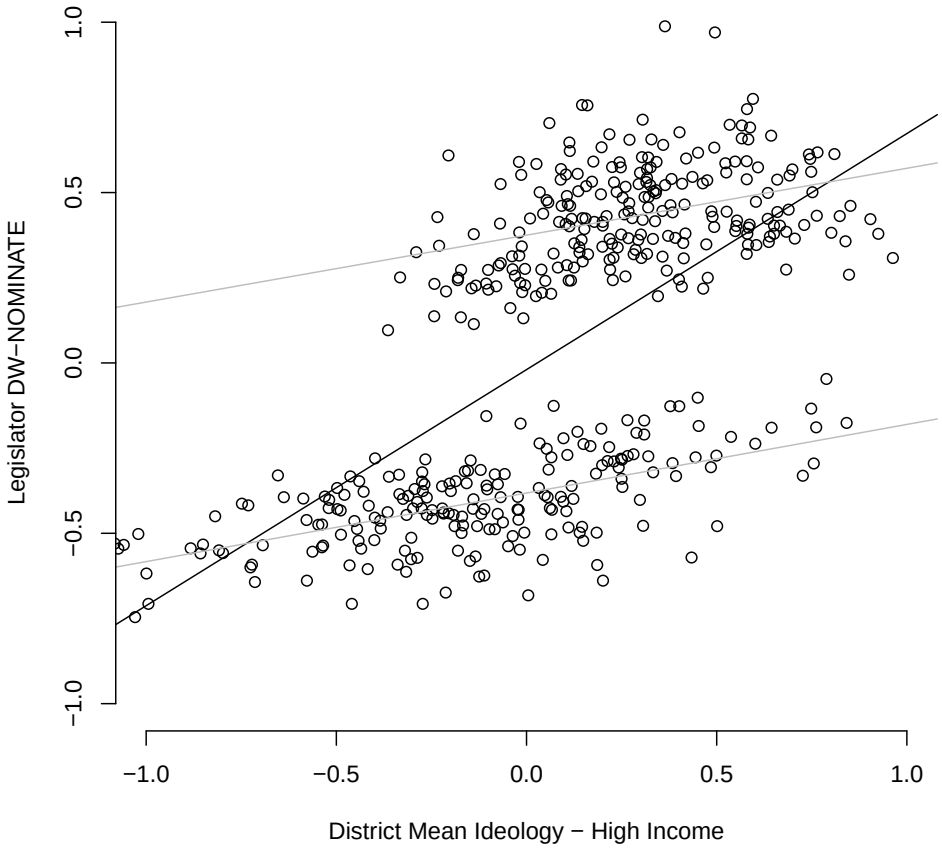

Figure 1: Relationship Between High Income Preferences and Legislator Position

of the lines are much steeper. At the same time, the reason for this steeper slope is quite apparent: there is much less variation in terms of positions. Low income voters are to the left of high income voters on average, but their estimated positions also tend to be closer to zero. We might expect to find this result if poorer constituents report their policy views with greater error.

Finally, Figure 3 shows what happens when we use the mean for the entire district to explain legislator positions. This variable explains more variance than either of the other two, with and R squared statistic of .51. And yet the key feature of the relationship remains:

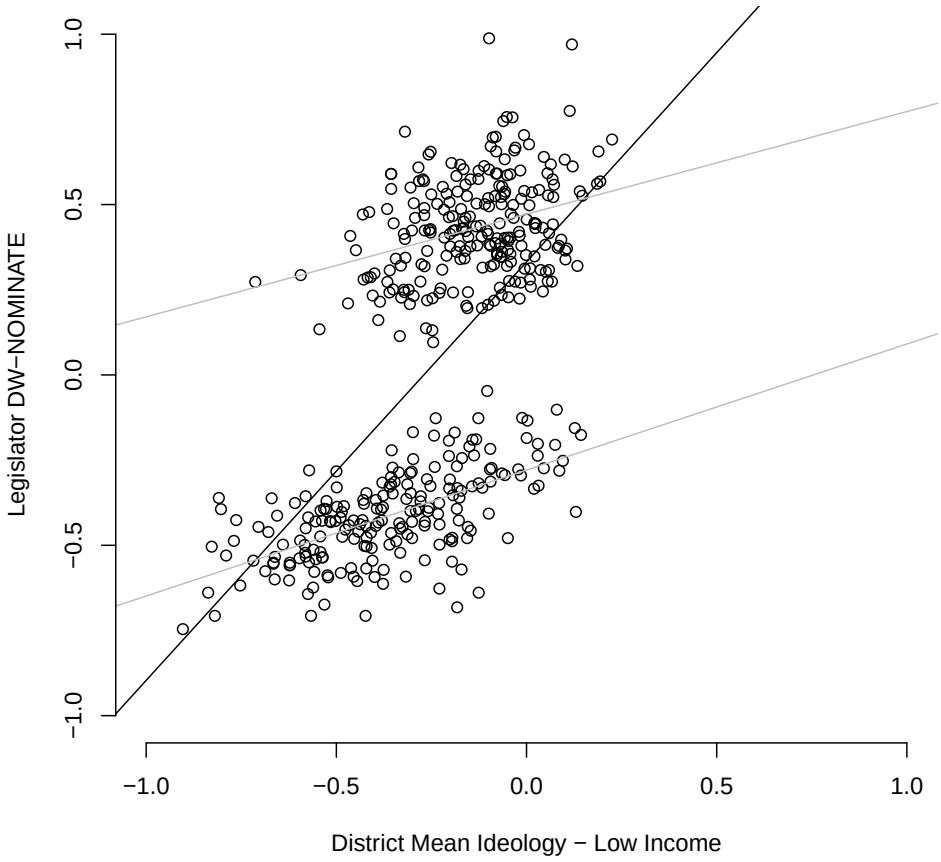

Figure 2: Relationship Between Low Income Preferences and Legislator Position

variance explained is mostly between-party (not too surprising, since the y variable has not changed), and our ability to explain within party variance is relatively poor. In fact these three Figures, 1, 2 and 3, are surprisingly similar. The differences in the relationship are overshadowed by the common disjuncture between the distribution of district opinion and the distribution of legislator positions.

One way to think about the substantive implications of these different relationships is to consider a hypothetical in which legislators were in fact responding only to low income constituents or only to high income constituents. We can estimate a univariate model within-

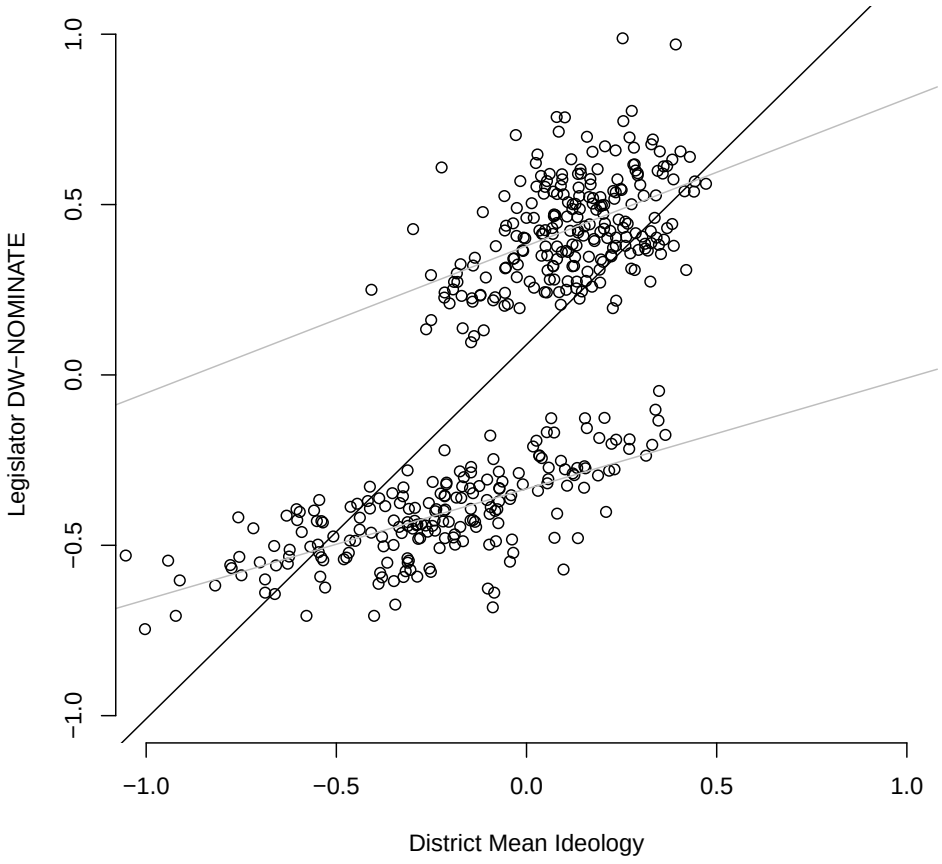

Figure 3: Relationship Between Mean District Preferences and Legislator Position

party for each group, generate predicted values, and examine which set of predicted values better matches reality. Figure 4 shows the result of this exercise, including the distribution of actual legislator positions. The distributions of predicted values in Figure 4 are less dispersed than the actual distribution of legislator positions. In contrast, they differ little from each other. This is evidence that noisy representation of all groups is much more significant than differences in representation between people at different income levels. The disparity between the two sets of predictions is hardly noticeable.

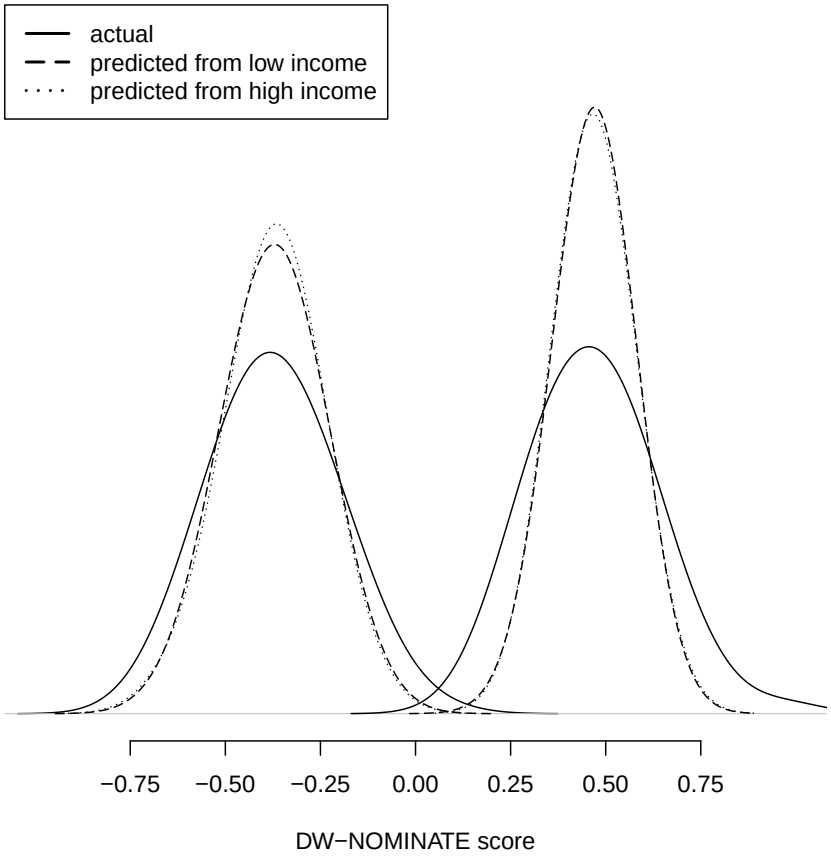

Figure 4: Predicted Positions for Representation of the Rich and the Poor

Results: Income and Legislator Partisanship

Figure 1, 2 and 3 show that our ability to explain within-party variation in legislator positions using constituent ideology is limited. As a result, we might think that a more reasonable model of representation is one where constituent ideology is responsible for the party of the representative but not their particular set of policy positions. For our purposes, the question becomes whether high income people are more important in deciding the party of the representative than low income people.

In order to test this hypothesis, I adapt the regression models from Table 3. Instead

of using a linear model where the dependent variable is the legislator's position, or DW-NOMINATE score, I employ a logistic regression model where the dependent variable is whether or not the legislator is a Democrat. Districts are more likely to be represented by Democrats when the population of those districts is more liberal. However, if higher income people have more importance in determining electoral outcomes, we might expect the preferences of higher income people to be a more important determinant of the partisanship of representatives.

Table 4 reports the result of a logistic regression model along the same lines as Table 3, but with party of the legislator as the dependent variable. The first two columns of the table show similar findings. In a model including only the preferences of low income constituents, the slope of the relationship between the mean preferences of low income constituents and the probability of electing a Democrat is significant. This relationship is also significant in a model where the mean preferences of high income constituents is the primary independent variable. The slope of this relationship is significantly less steep, although the model fits somewhat better. In both cases, the relationship is negative, as expected: more conservative constituencies are less likely to elect Democrats.

The third column of Table 4 diverges somewhat from previous results. Although the relationship between the preferences of each group and the probability of electing a Democrat all have the expected sign, only two of the coefficients are significant. The positions of high-income people have by far the largest slope and the relationship does not steadily decrease with income as it did before. The "medium-high" group appears to have the smallest slope, while the low income and medium-low income groups are in a close tie for second place. The coefficent on the preferences of the rich is significantly greater than the rest, while

Table 4: Logistic regressions explaining legislator party

	Dependent variable:		
	Democratic Legislator		
	(1)	(2)	(3)
Constant	-2.36	0.60	-0.43
	(2.660)	(2.758)	(3.182)
μ_L	-7.72^{***}		
	(0.845)		
μ_H		-5.28^{***}	
		(0.544)	
$p_H \times \mu_H$			-20.42^{***}
			(4.878)
$p_{MH} \times \mu_{MH}$			-4.34
			(3.137)
$p_{ML} \times \mu_{ML}$			-9.25^{**}
			(4.195)
$p_L \times \mu_L$			-9.01
			(5.604)
p_H	-1.47	-2.83	-3.99
	(4.686)	(4.881)	(5.696)
p_L	14.11^{***}	15.78^{***}	11.77^{**}
	(4.012)	(4.259)	(4.581)
p_{ML}	-5.29	-7.54	-2.63
	(6.165)	(6.534)	(7.269)
Observations	435	435	435
Log Likelihood	-213.139	-196.190	-174.659
Akaike Inf. Crit.	436.278	402.379	365.319
Note:			*p<0.1; **p<0.05; ***p<0.01

the coefficient on the preferences of the other groups are all statistically indistinguishable. Notably, in all specifications the proportion of the district that is poor has a substantial positive effect on the probability of electing a Democrat.

Regression coefficients are notoriously difficult to interpret in logistic regression models. In order to understand the substantive difference between the effects for the rich and the poor, a visualization is helpful. I use a model very similar to the one from the last column of Table 4 to show the predicted change in the probability of electing a Democrat based on a change in the ideology of the mean low and high income constituent, respectively. The only difference between the model used to calculate these probabilites and the model in Table 4 is that it includes all direct effects, following the folk wisdom on using interactions in regression models (Brambor, Clark, and Golder, 2006). This change in specification makes very little difference in the resulting probabilities. To calculate the change in probability I hold all other variables besides the variable of interest at their mean.

Figure 5 shows the result of this exercise. The left panel shows the predicted probability of a Democratic legislator given the preferences of low income constituents, and the right panel shows the predicted probability of a Democratic legislator given the preferences of high income constituents. As expected from the regression table, the slope is steeper for high income constituents. Nonetheless, the slope in the left panel is not flat: Democrat legislators are substantially more likely when poor constituents are liberal than when they are conservative.

The "rug" for each graph shows the distribution of the x values. As discussed, variation in preferences is substantially greater for high income than low-income constituents. In this case, the range where the slope for high income constituents exceeds the slope for low

Figure 5: Probability of a Democratic legislator by ideology of income groups

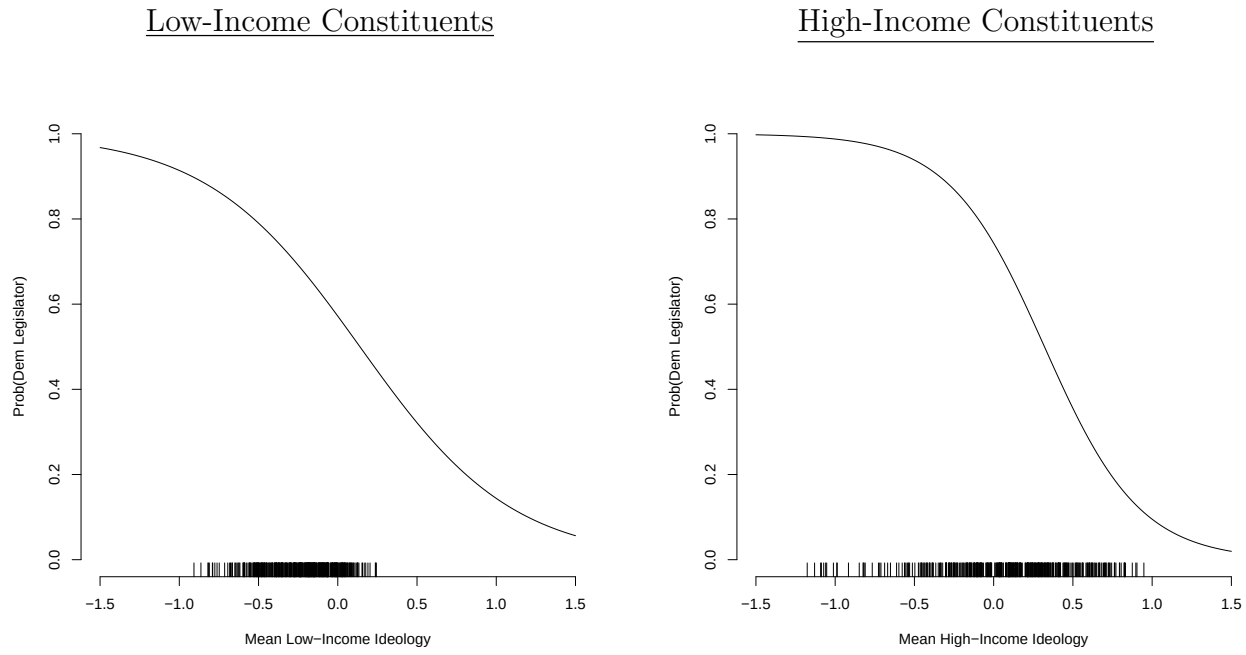

The y-axis is the predicted probability of having a Representative from the Democratic Party. This comes from a model very similar to the one in colum three of Table 4, but with all direct effects included. Probabilities are calculated with all variables except the one on the x-axis (mean low-income ideology and mean high-income ideology, respectively) held fixed at their mean. The tick marks on the bottom of the graph show the distribution of the actual x values.

income constituents occurs in a region where there is no data for low income constituents. Specifically, there are very few districts where the mean ideology of low income constituents is to the right of zero. In contrast, for high income constituents, most of the decline in the likelihood of electing a Democrat occurs to the right of zero. Our analysis has told us the slope is steeper in the panel right, but this steep slope occurs in a region where we have no data in the left panel. In other words, we cannot tell how unlikely electing a Democrat would be in a district with very conservative poor residents, because no such district exists.

Conclusion

In this paper I have shown that while there is some evidence that lower income people are less represented, this finding lacks robustness and the substantive importance of this difference is limited in comparison to the gap in overall representation. Given the small extent of differential representation and uncertainty regarding it, priority should be given to understanding overall representation and why the distribution of legislator preferences is so different from the distribution of average preferences in districts. Separating the public into large categories by income does not seem to help us solve this puzzle.

In order to understand how representation is unequal, better theories of representation are needed. A good theory of representation should account for the fact that most of the variance in legislative positions is within party. With such a theory, political scientists would be better able to evaluate whether legislators take the preferences of their constituents into account without regards to income. In this paper I have used two very basic theories: that legislators represent the mean voter, and that the party of the legislator is determined by

the mean voter. Future research should strive to build richer theories of representation from microfoundations.

A likely explanation for the small amount of differential representation that I find is that there is more measurement error in the preferences of the poor than the rich. This affects the analyis insofar as preferences have been imperfectly observed. However, legislators as well as political scientists may have more difficulty observing these preferences. Future research should investigate this possibility as a likely cause of a weak representational link.

One argument that I have not made in this paper is to say that low income people are in fact well represented. Work that makes the argument that the political system does not represent the poor very well may be right. Much depends on a value judgement about what aspect of preferences should be represented. Certainly the political system has not responded to the economic needs of lower income people in a way that standard political economy models would predict (Bonica et al., 2013; Hacker and Pierson, 2011). The findings of this paper may plausibly answer the question, as Bonica et al. (2013) puts it, of "why hasn't democracy slowed rising inequality?" If legislators respond only weakly to their constituents in general, and *perhaps* their low income constituents in particular, then not one but two conditions are violated that would be needed to beget a truly *democratic* response to rising inequality. The question of what institutions might improve both of these links, and the downstream effects of doing so, is an urgent matter for future research.

References

Achen, Christopher H, and Larry M Bartels. 2012. "Blind Retrospection: Why Shark Attacks Are Bad For Democracy." *Center for the Study of Democratic Institutions, Vanderbilt University. Working Paper* .

Ansolabehere, Stephen, Jonathan Rodden, and James M Snyder. 2008. "The strength of issues: Using multiple measures to gauge preference stability, ideological constraint, and issue voting." *American Political Science Review* 102(02): 215–232.

Bartels, Larry M. 2009. *Unequal democracy: The political economy of the new gilded age.* Princeton University Press.

Black, Duncan. 1948. "On the rationale of group decision-making." *The Journal of Political Economy* pp. 23–34.

Bonica, Adam, Nolan McCarty, Keith T Poole, and Howard Rosenthal. 2013. "Why Hasn't Democracy Slowed Rising Inequality?" *The Journal of Economic Perspectives* 27(3): 103–123.

Brambor, Thomas, William Roberts Clark, and Matt Golder. 2006. "Understanding interaction models: Improving empirical analyses." *Political analysis* 14(1): 63–82.

Brunner, Eric, Stephen L Ross, and Ebonya Washington. 2013. "Does Less Income Mean Less Representation?" *American Economic Journal: Economic Policy* 5(2): 53–76.

Butler, Daniel M. 2014. *Representing the Advantaged: How Politicians Reinforce Inequality.* Cambridge University Press.

Campbell, Angus, Philip E Converse, Warren E Miller, and E Stokes Donald. 1966. *The American Voter*. New York: Wiley.

Caplin, Andrew, and Barry Nalebuff. 1991. "Aggregation and social choice: a mean voter theorem." *Econometrica: Journal of the Econometric Society* pp. 1–23.

Clinton, Joshua D. 2006. "Representation in Congress: constituents and roll calls in the 106th House." *Journal of Politics* 68(2): 397–409.

Clinton, Joshua, Simon Jackman, and Douglas Rivers. 2004. "The statistical analysis of roll call data." *American Political Science Review* 98(02): 355–370.

Downs, Anthony. 1957. "An economic theory of political action in a democracy." *The journal of political economy* pp. 135–150.

Erikson, Robert S, and Yosef Bhatti. 2011. "How Poorly Are the Poor Represented in the US Senate?" *Who gets represented* pp. 223–46.

Fiorina, Morris P, and Samuel J Abrams. 2012. *Disconnect: the breakdown of representation in American politics*. Vol. 11 University of Oklahoma Press.

Gilens, Martin. 2012. *Affluence and influence: Economic inequality and political power in America*. Princeton University Press.

Gilens, Martin, and Benjamin I Page. 2014. "Testing theories of American politics: Elites, interest groups, and average citizens." *Perspectives on Politics* .

Hacker, Jacob S, and Paul Pierson. 2011. *Winner-Take-All Politics: How Washington Made the Rich Richer–and Turned Its Back on the Middle Class*. Simon and Schuster.

Hill, Seth J., and Chris Tausanovitch. 2014. "Non-Polarization in the American Public." *Manuscript* .

Poole, Keith T, and Howard Rosenthal. 1997. *Congress: A political-economic history of roll call voting*. Oxford University Press.

Schofield, Norman. 2007. "The mean voter theorem: necessary and sufficient conditions for convergent equilibrium." *The Review of Economic Studies* 74(3): 965–980.

Tausanovitch, Chris, and Christopher Warshaw. 2013. "Measuring Constituent Policy Preferences in Congress, State Legislatures, and Cities." *The Journal of Politics* 75(02): 330–342.

Tausanovitch, Chris, and Christopher Warshaw. 2014. "Do Legislator Positions Affect Constituent Voting Decisions in U.S. House Elections?" *Manuscript* .

Zaller, John. 1992. *The nature and origins of mass opinion*. Cambridge university press.