

# Measuring the Intensity of Policy Preferences \*

Chris Tausanovitch<sup>†</sup>  
Department of Political Science  
UCLA

July 2018

*Preliminary draft.*

This paper proposes a new method for measuring the relative intensity of policy preferences. It is motivated by conjoint analysis but unlike conjoint analysis the measure here is meant to measure intensity at the individual level. Having individual level measures is necessary for evaluating existing measures. Individual measures are also needed to evaluate the primary hypothesis of the issue publics literature: that intensity is highly concentrated. Attitude importance measures do a poor job of capturing preference intensity. The proposed measure shows that intensity is indeed highly concentrated, as the literature suggests. Variation in preference intensity may help explain anti-majoritarian policymaking and unequal representation.

---

\*For funding a project which ultimately led to this one, I thank the UCLA Academic Senate. For feedback on the survey instrument I thank Stephanie Carrie, Michael Castle Miller, Joanna Castle Miller, Alex Tausanovitch and Zachary Tausanovitch. I thank Seth Hill, Greg Huber, Jeff Lewis, Jason Oh, Alex Tausanovitch and Chris Warshaw for helpful comments.

<sup>†</sup>Assistant Professor, Department of Political Science, UCLA, [ctausanovitch@ucla.edu](mailto:ctausanovitch@ucla.edu).

# 1 Introduction

The preference intensity of voters over policies is of clear importance for central questions in political science. Differences in intensity can create wide disparities in representation of different groups (Fiorina, 1974). We can reasonably expect stricter policy accountability when voters have more intense preferences. Intensity is also a key variable for understanding whether people will pay the costs of voting at all (Riker and Ordeshook, 1968).

According to the formal theory literature, preference intensity is the willingness to incur a sacrifice to achieve a given end (e.g. Arrow, 1951, p. 114). This notion is related to the notions of “issue publics” or “attitude strength” in the psychological and empirical literature, but not quite the same. According to Converse (1964) an issue public is the small set of people who have crystallized views on an issue, and therefore “care” more about it. This is often measured by membership in groups that have a lot at stake in a given policy area (e.g. Bishin, 2009). “Attitude strength” characterizes views that are stable, consequential, and difficult to change (Krosnick and Petty, 1995), often measured by a standard question about issue importance (e.g. Krosnick, 1990). I will show that this approach does a poor job of capturing intensity as it is measured here.

In this paper I propose a measure that is more directly linked to the theoretical concept of preference intensity. I begin by asking respondents to rate their subjective preference for several different policy options within three different issue areas. Then I ask respondents to rate candidates who take one of these positions on each of the three policies. For every respondent, I model their candidate rating as a function of their ratings for each of the candidate’s policy positions, with coefficients varying by issue area. Intensity is measured by the difference in ratings that results from the candidate’s choice of one policy over another once the estimated coefficients are applied. This captures the respondents’ willingness to substitute policies for one another or for other candidate attributes- the degree of sacrifice they are willing to incur.

I validate the measure using conjoint analysis. Conjoint analysis can be used to identify

the group-level treatment effect of a given candidate position on candidate ratings (Hainmueller, Hopkins, and Yamamoto, 2013). I show that when the model-based measures for individuals are aggregated, they are strongly related to the conjoint estimates for groups.

The most common measure of attitude strength does a poor job of capturing intensity. However, the measure of intensity that is used here confirms the claim that intensity is highly concentrated (Converse, 1964; Krosnick, 1990). The most intense decile typically puts more than three times as much weight on a given policy choice than the fifth decile does. This raises the prospect that differences in intensity can lead to differences in representation. If voting is stochastic then candidate positions have a greater impact on the vote probabilities of voters with intense preferences<sup>1</sup>. In other words, candidates have an incentive to give much greater weight to the voters with the most intense preferences.

None of the conclusions of this paper can be automatically generalized to other issue areas. External validity remains a serious hurdle, but for this paper I regard the idealized set up as a feature, not a bug. My argument is that variation in preference intensity is poorly understood relative to its importance in theoretical and normative accounts of democracy. The goal of this paper is to advance this avenue for the understanding of policy representation.

## 2 Methodology

The best existing method for measuring preference intensity is conjoint analysis<sup>2</sup>. Conjoint analysis allows us to measure the effect of policy positions on support for candidates when these policy positions are weighed against one another and other factors. It does so through the gold standard of randomized intervention. If all attributes are assigned at random, then support for candidates can be compared among these various treatment groups.

---

<sup>1</sup>More precisely, if a voter is intense with respect to a particular choice between two policy options, then the selection of one or the other can have a large impact on that voter's probability of selecting the candidate in question. Voters who are have intense preferences with regard to a particular choice may not have intense preferences with regard to every choice between two policies in some narrowly defined policy area.

<sup>2</sup>Some examples include Ballard-Rosa, Martin, and Scheve (2017), Bechtel and Scheve (2013) and Tomz, Weeks, and Yarhi-Milo (2018).

However for the question at hand, conjoint analysis suffers from a serious drawback. Like standard single-treatment experiments, conjoint analysis can only estimate aggregate effects. This means it can be used to measure the relative intensity of the population with respect to one choice or another, but not the intensity of a single individual <sup>3</sup>. With large enough samples we may use conjoint analysis to compare effects across pre-defined groups, but we can't tell which groups will be the most intense a priori.

In order to surmount this obstacle I use an observational measure that is strongly linked to the theoretical concept of interest. I first ask respondents to rate a set of policies, giving higher ratings to the ones they prefer more. These policies are organized into policy areas. Then I ask them to rate a set of candidates, whose policy positions are drawn from the policy questions. Candidates take one position in each policy area, drawn at random. I also ask respondents to rate the two major political parties, and randomly assign each candidate to a party. Then for each respondent I model the candidate rating as a function of the ratings of their policy positions and party, with coefficients that are allowed to vary by policy area. I am interested in the effect of different policy positions on the predicted values from this regression. The greater impact a given change in position has on these predicted values, the more intense the respondent is with respect to that change in position.

Consider a set of candidates indexed by  $i$ , where  $imm_i$  is their position on immigration,  $tax_i$  is their position on taxes and  $gun_i$  is their position on guns. The respondent has provided a set of ratings for every policy that will be considered in each of these issue areas.  $imm\_rating_{imm_i}$  is the respondent's rating of candidate  $i$ 's position on immigration,  $tax\_rating_{tax_i}$  is the respondent's rating of candidate  $i$ 's position on taxes, and  $gun\_rating_{gun_i}$  is the respondent's rating of candidate  $i$ 's position on guns.  $party_i$  is the candidate's political party, and  $party\_rating_{party_i}$  is the respondent's rating of that political party.  $rating_i$  is the respondent's rating of candidate  $i$ . I model this as follows:

---

<sup>3</sup>This is a claim about feasibility. Conjoint analysis could be used to estimate intensity for an individual under the assumption that repeated experiments are independent. However it would take an extremely large set of experiments to obtain a precise estimate. In practice there are limits to the number of experiments that can be conducted on a single individual.

$$rating_i = \beta_0 + \beta_1 imm\_rating_{imm_i} + \beta_2 tax\_rating_{tax_i} + \beta_3 gun\_rating_{gun_i} + \beta_4 party\_rating_{party_i} + \epsilon_i \quad (1)$$

Where  $\epsilon_i$  is a normally distributed independent disturbance.  $\beta_0$  is the intercept and  $\beta_1$ ,  $\beta_2$  and  $\beta_3$  are the coefficients for immigration, taxes, and guns respectively.  $\beta_4$  is the coefficient for party. This regression is run for each respondent, so the coefficients are all respondent-specific.

The intensity of a choice between two policies is measured by the difference in predicted values from this regression. So for instance, a respondents' intensity with respect to the choice between immigration policy 1 and immigration policy 2 is the difference between the predicted values from equation 1 when  $imm\_rating_{immigration\ policy\ 2}$  is substituted for  $imm\_rating_{immigration\ policy\ 1}$  when all other variables are held at a fixed value. In this case, the estimated difference is simply  $\beta_1(imm\_rating_{immigration\ policy\ 2} - imm\_rating_{immigration\ policy\ 1})$ . More generally, take any set of profiles A and B, where each profile entails an immigration policy, a tax policy, a gun policy and a party. Then the estimated intensity with respect to this choice is given by:

$$\begin{aligned} \widehat{intensity}_{A,B} = & \\ & (\hat{\beta}_1 imm\_rating_{imm_A} + \hat{\beta}_2 tax\_rating_{tax_A} + \hat{\beta}_3 gun\_rating_{gun_A} + \hat{\beta}_4 party\_rating_{party_A}) \\ & - (\hat{\beta}_1 imm\_rating_{imm_B} + \hat{\beta}_2 tax\_rating_{tax_B} + \hat{\beta}_3 gun\_rating_{gun_B} + \hat{\beta}_4 party\_rating_{party_B}) \end{aligned} \quad (2)$$

It may be useful in some circumstances to think about intensity with regard to a multifaceted choice. For instance, we might wish to approximate the intensity of views about actual candidates. However in this paper I will focus on the intensity of choices within policy areas and how these compare.

This measure of intensity may seem odd from the perspective of existing measures because the unit of observation is a choice, not an issue area. It does not measure how much someone cares about “guns” but how much they care about the choice between two specific gun policies. This is motivated by the fact, shown below, that voters who care very much about one policy choice in a given issue area may care very little about another policy choice in that area, even if the latter choice concerns a “big” difference in facts on the ground. For instance, I will show that many voters who intensely want taxes on high income individuals to be raised also do not care much whether they are raised by a moderate amount or a very large amount. Measures of how much people care about an issue mask important variation in how intense they are with respect to particular alternatives.

Nonetheless, in some circumstances it may be useful to have a single-number summary that averages the intensity of all of the relevant comparisons within an issue area. One such measure is the mean absolute intensity of all adjacent comparisons. This has the catchy acronym MAIAAC. This measure only makes sense for a policy area where the options have a natural ordering. For a policy area where the options do not have a natural ordering, the mean absolute intensity of all comparisons (one fewer A: MAIAC) could be used, assuming that the number of policies makes this computationally tractable.

### **3 Data**

I measure preference intensity with respect to three policy areas, with seven distinct policies in each area. These policies are shown in Table 1. The first set of policies concern the treatment of undocumented immigrants. As a convenient shorthand I will refer to this policy area by the over-general term “Immigration.” The policies range from providing immediate citizenship to all undocumented immigrants to immediately deporting all undocumented immigrants. In each policy area the policies are numbered 1 to 7 with 1 being the most liberal policy and 7 the most conservative. The policies are meant to vary along as few

features as possible so that there is no ambiguity about which policy is more conservative. However this is balanced against a desire to capture variation across the political spectrum and to capture some of the options actually discussed. For example, in the immigration area the most liberal policy offers the most and easiest legal inclusion to the greatest number of undocumented immigrants, and the most conservative policy offers the least. However policies 4 and 5 limit legal inclusion to undocumented immigrants who were brought to the United States as children. This introduces a different consideration that could lead to ratings that violate single-peakedness. In the results section I will show that this is true in some individual cases but not on average.

The second policy concerns taxes on high income individuals (“Taxes”). They range from taxing income over the first \$600,000 (the status quo top tax bracket) at 100% to reducing the tax on income over the first \$600,000 to 20%. The third policy area concerns limitations on firearms, ranging from simply “Ban all guns” to “Allow all Americans to own any gun without restrictions.” The intermediate categories focus on limitations on “assault weapons.”

My approach begins by asking respondents to rate how much they like a given policy on a scale from 0 to 100. They are told that 0 is reserved for a policy they “truly hate” and 100 is for a policy they “truly love.” The ratings scale is also labelled at seven equidistant points, ranging from “truly hate” to “hate,” “dislike,” “neutral,” “like,” “love,” and “truly love.” This question is meant to elicit each respondent’s subjective preference for a given policy. The endpoint labels are meant to prevent floor and ceiling effects by emphasizing extreme utility and disutility. Respondents are shown all of the rating questions for all three policy areas on a single page, and they are told to try to make these ratings comparable to the other ratings across the entire section. Finally, a bold, underlined statement tells them to “Only give two policies the same score if you would trade one for the other.”

On the same page I also ask respondents to rate the two major U.S. political parties on the same scale. This is meant to create a point of comparison as well as allowing respondents to trade off policy with at least one important non-policy factor when it comes to making

**Table 1:** Policies. Respondents are asked to choose their most preferred policy in each of these three policy areas. Then they are asked to rate each policy in each area. Finally, they are asked to rate and choose between candidates that take one stance in each area.

Treatment of undocumented immigrants (“Immigration”)	
1	Provide immediate citizenship to all undocumented immigrants who have no criminal record
2	Create a path to citizenship for all undocumented immigrants who have no criminal record
3	Create a path to citizenship only for undocumented immigrants who are employed, English-speaking, have no criminal record and have been in the US for many years
4	Create a path to citizenship only for undocumented immigrants who were brought to this country as young children
5	Create a legal status but no path to citizenship for undocumented immigrants who were brought to this country as young children
6	Do not create a legal status for any undocumented immigrants but exercise discretion in deportations
7	Deport all undocumented immigrants immediately
Taxes on high income individuals (“Taxes”)	
1	Tax income over the first \$600,000 at 100%
2	Tax income over the first \$600,000 at 80%
3	Tax income over the first \$600,000 at 60%
4	Tax income over the first \$600,000 at 50%
5	Keep the tax on income over the first \$600,000 at the current rate of 37%
6	Reduce the tax on income over the first \$600,000 to 30%
7	Reduce the tax on income over the first \$600,000 to 20%
Regulation of ownership of firearms (“Guns”)	
1	Ban all guns
2	Ban all guns except single-shot firearms
3	Ban assault rifles and limit firearm magazines to 10 rounds or less
4	Ban assault rifles
5	Require mandatory training and enhanced background checks for purchase of assault rifles
6	Do not restrict ownership of assault rifles beyond current law
7	Allow all Americans to own any gun without restrictions

judgements about candidates.

I ask respondents to rate twenty hypothetical candidates on the same 0 to 100 scale used above. Each candidate takes a randomly assigned position on each issue as well as having a randomly assigned party. To mitigate concerns that respondents may think in terms of partisan strategy or control of Congress, they are told to assume that each candidate would enact all of his or her preferred policies if elected. They are told “Only give two candidates the same score if you would trade one for the other” and shown ten candidates per page.

The data for this project comes from a survey run from April 4 to 11, 2018 using the Luc.id Fulcrum Exchange, an online marketplace for recruiting survey respondents. Fulcrum provides sample in real time to meet a set of demographic quotas that match census demographics. It draws from a pool of more than 350,000 unique respondents daily. Coppock and McClellan (2017) show that Fulcrum outperforms Amazon Mechanical Turk in terms of its ability to match estimates from probability samples and replicate well-established findings



in political science and psychology<sup>4</sup>.

In Appendix I show the survey demographics and compare them to other sources. The demographics are quite similar to those from the 2016 American National Election Study. On the survey I asked a question about whether America is on the “right track.” In the Appendix I compare the demographic cross tabulations on that question with a cross tabulation on the same question from a contemporaneous iteration of the well-known Economist/YouGov survey. The results are again similar. For simplicity I dispense with survey weights for the analyses that follow, given that the demographics are fairly close to the general population and that external validity looms larger as a concern here than the (legitimate) concern about representativeness.


In addition to asking the ratings questions and the candidate choice questions I ask the traditional attitude strength question <sup>5</sup> and a simple “most preferred policy” question for each issue area as well as party identification and a right track/wrong track question (for the comparison to the YouGov survey in the Appendix). I also ask respondents ten candidate choice questions. Unbeknownst to the respondents, these are drawn from candidates that they also rate. Half are randomized to have a “neither” option.

Analyses do not include respondents who completed the survey in less than 5 minutes, or who completed it in under 9 minutes and failed the screener question, about 13% of the sample. The total sample size after these cases are dropped is 1,621 respondents. Median completion time for the survey was 14 minutes. Another 5% of the sample was dropped for leaving more than 3 blanks on the rating questions or using less than 4 unique values to answer these questions. As I will show, this leaves many responses that do not appear very credible, however I have left them in the sample in an attempt to avoid prejudging the

---

<sup>4</sup>It should be noted that unlike conventional survey research firms, the Fulcrum Exchange gives the researcher considerable control over the fielding of surveys. The researcher controls the timing, the pricing, and the quotas, and may adjust these in real time as the responses come in. Researchers can thereby affect the quality of the results, for better or for worse.

<sup>5</sup>Specifically, I ask respondents to “Please indicate how important each of the following policy areas are to you personally:” The areas are called “Treatment of undocumented/illegal immigrants,” “Gun control/gun rights,” and “Taxes on high income people.” As is common practice, there are five responses ranging from “Not at all important” to “Extremely important.”



**Figure 1:** Most preferred policies. Each bar represents the set of respondents that selected the policies, 1-7, in each policy area. The light grey parts of the bars represent Democrats, the dark grey parts Independents and the black parts Republicans. The specific policies are listed in Table 1

conclusions. For instance, it is not possible a priori to say whether a respondent gave random values for the rating questions or whether they simply have unusual preferences.

Figure 1 shows responses to the initial “most preferred policy” questions for each policy area, stacked by party ID. The light grey parts of the bars represent Democratic respondents, the dark grey parts of the bars represent pure Independents and the black parts of the bars represent Republicans. In each case, some set of moderate categories are selected at much higher rates than the most extreme categories. As we would expect, Republicans select the more conservative policies (to the right) at much higher rates than Democrats do, and likewise eschew the more liberal policies.

## 4 Results

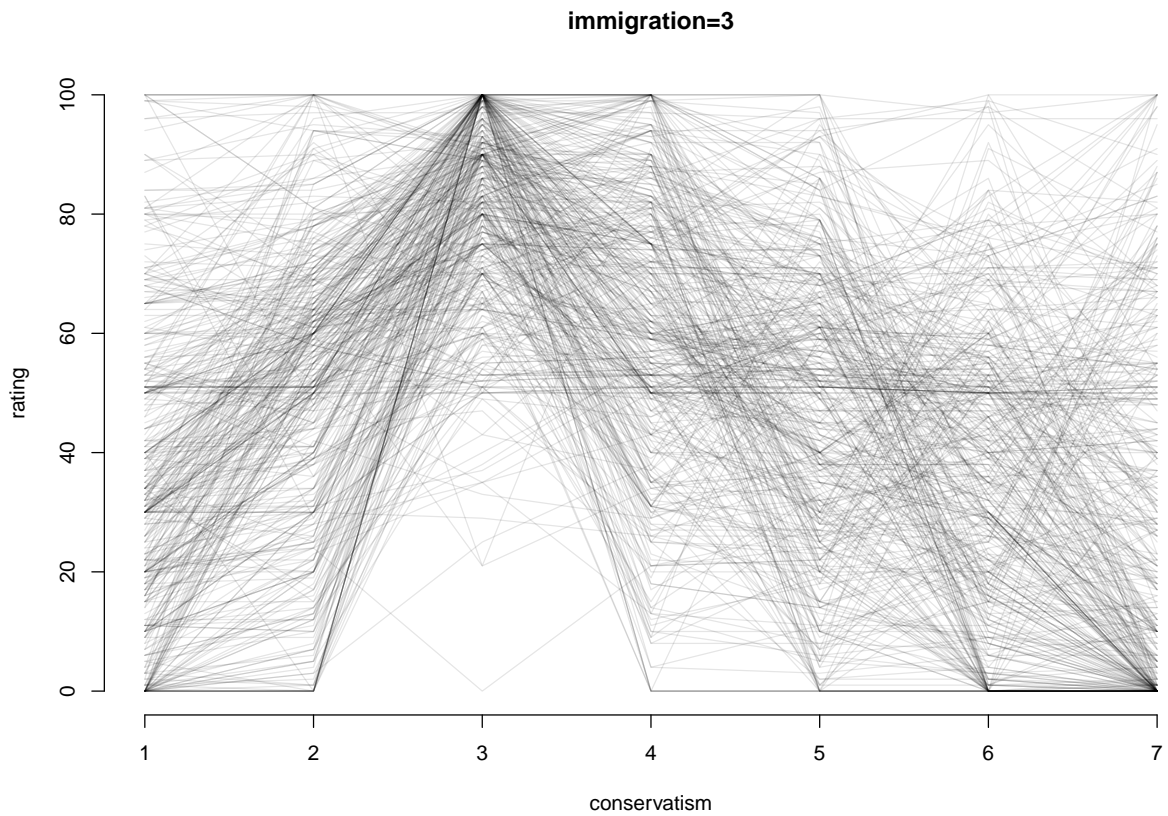
Figure 1 captures only a narrow slice of information about each respondent’s preferences: their *most* preferred policy. It tells us little about *how much* they prefer this option over others, whether their preferences are single-peaked, what is their least preferred option, and so on. Figure 2 takes only the set of respondents who chose policy 3 on the immigration

question (a path to citizenship with restrictions), and graphs each respondent’s rating for every category, connected by a line. This can be thought of as a “pseudo utility function” in the sense of “decision utility” (Kahneman, Wakker, and Sarin, 1997).

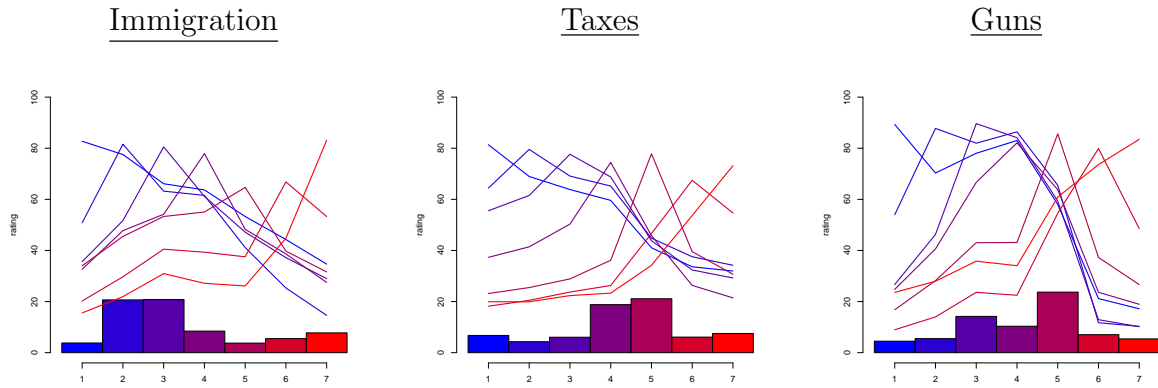
Figure 2 makes clear that there is substantial heterogeneity in ratings among people who most prefer option three. Most respondents give a high rating for this option, represented by the dense set of triangular shapes at point three, with very few points below 50. However there is a very high degree of variation in how steeply these ratings drop off towards options 2 and 4. Some of these connecting lines are very steep, going almost to 0. Others are quite shallow, representing respondents who don’t see much of a difference between these options.

Figure 2 also makes it clear that there are respondents whose ratings are starkly contrary to our expectations. We can see that there are ratings that are not single-peaked, and in some cases do not even peak at 3, the supposed most preferred option. There are a number of cases of ratings that are close to being flat, often near 50. Some respondents give very high ratings to policies 1 and 7, indicating a high degree of bimodality. Some of these responses are strange to the point of implausibility, but to remove them would be a form of selecting on the dependent variable. Others may represent exactly what we are looking for. If a respondent does not particularly care about this issue, then their ratings should be close to flat. Appendix B shows the complete set of ratings for every response category in each policy area.

While the individual ratings may seem chaotic, the means are much more well-behaved. Figure 3 graphs the average raw ratings for each policy area. Each line represents the average ratings for a set of respondents who chose a given policy as their most preferred, with colors corresponding to the histogram of responses shown at the bottom of the graph (these correspond to Figure 1). The most liberal categories are represented as blue, the most conservative categories red, with shades in between for intermediate categories. Each line peaks at the correct point, and with the exception of some points of almost-indifference, each is single peaked. One can easily tell what category corresponds to each line simply by



**Figure 2:** Individual ratings. The y-axis is the rating given, and each line represents a respondent. The x-axis is the immigration policy being rated, with more conservative policies to the right. All responses shown are for people who said that policy 3 was the closest to their own view in a previous question.



**Figure 3:** Average policy ratings. Each line represents the average ratings of each policy by a set of respondents. The blue line represents those who said the most liberal policy was the closest to their own view. The red line represents those who said the most conservative policy was the closest to their own view. Intermediate shades represent the five intermediate groups. These are reflected in the colors of the corresponding histogram, which is proportional to the number of respondents who selected each option as closest to their own view.

examining the shape of the line and where it peaks.

If taken at face value, this graph illuminates some interesting patterns. The range of ratings is not much greater for extreme respondents than for more moderate respondents. The slopes of the lines vary over different ranges in ways that are telling. For instance, extreme liberals are relatively indifferent between extremely liberal categories and moderate liberal categories, particularly when it comes to taxes and guns. Extreme conservatives are just the opposite, with most of the change in their utility functions taking place between relatively conservative categories. Moderates typically give quite low ratings to policies that are extreme.

One pattern that is more doubtful is the similar variation in ratings across issue areas. Despite being implored to make ratings comparable across each area, respondents may have found it difficult to compare so many different policies. Respondents may feel that it is socially desirable to label each area as important. Moreover, they may lack an understanding of their own preferences. Asking a person to weigh their feelings about a policy in the abstract is different than asking them to weigh them in an actual decision.


A better approach is to actually force respondents to make tradeoffs. This is what the candidate rating questions do. The policy ratings above are the inputs into equation 1.

In order to validate the estimates from this model, I examine the change in candidate rating that would be produced if a candidate were to move from policy 1 in a given area, the most liberal policy, to policy  $X \in \{2...7\}$ , according to the model. For each policy option, I take all of the respondents who most preferred that option and I calculate six different average effects for that group, one for each  $X$ . For each of the three policy areas, this means calculating six average effects for seven groups of respondents, for a total of 126 effects. The full set of effects is shown in Appendix C.


Figure 4 compares the effects estimated via conjoint analysis with the effects estimated via the model (right panel). The effects from the model explain 81% of the variation in the conjoint effects. The average magnitude of these effects are 67% of the size of the conjoint effects. In other words, we can account for two thirds of the conjoint effects using this method.

The fact that these effects are strongly linked to the conjoint effects provides aggregate-level evidence of their validity. Figure 5 plots the average predicted values from the model, following the approach from Figure 3. In order to produce predicted values I assume that all candidates are Democrats candidates and that the candidates take on the average rating for the auxiliary policies. I calculate the predicted values when the policy of interest is varied, and graph the results for each group. The assumed baseline affects the levels of the prediction but not the shape because the underlying model is linear.

Most of the patterns from Figure 3 hold in Figure 5, but the overall range is compressed. This reflects the fact that these ratings entail trade-offs between different candidate attributes. Extreme liberals continue to appear relatively indifferent between extremely liberal categories and moderate liberal categories and extreme conservatives appear most intense with respect to conservative categories. Once again, extreme voters are not much more intense than other voters. Overall the range of ratings by respondents who answer in categories



**Figure 4:** Comparison of modeled ratings to conjoint results. Each point corresponds to the set of respondents who chose a particular most preferred policy in a given issue area and the estimated effect of a given candidate policy position on their candidate rating. These effects are calculated relative to a baseline where the candidate takes position 1, the most liberal position. The y-axis is the effect calculated using conjoint analysis and the x-axis is the effect calculated using the average predicted values from a set of linear regressions for each respondent, where the independent variables are the respondents' ratings of candidate policies and party. The dashed line is the 45-degree line and the black line is the regression line. Blue dots are effects on gun policy, red dots are effects on immigration, and green dots are effects on taxes.



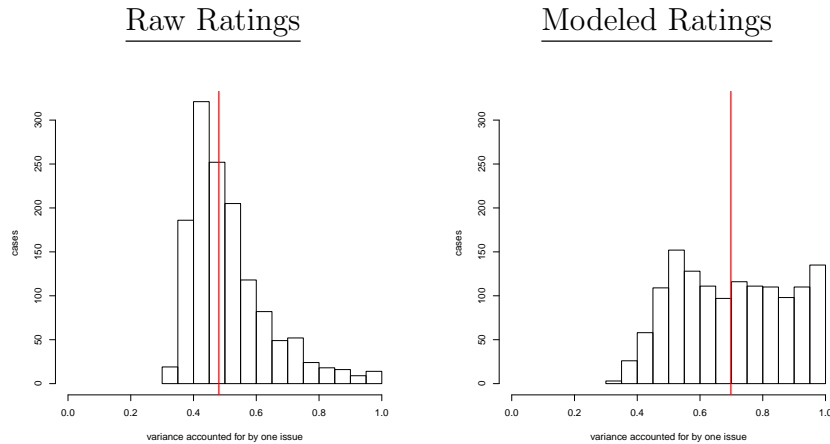
**Figure 5:** Average modeled ratings. As in 3, each line represents a set of respondents who chose a particular most preferred policy in each issue area, from liberal (blue) to conservative (red). In this figure the y-axis represents the predicted value for the candidate ratings when the policy is varied, with other policy ratings set at their mean. It is assumed that the candidate is a Democrat. Higher numbers on the x axis represent more conservative candidate positions.

1 or 7 is 21.3 on average, with a median of 14.6. The range of ratings by respondents who answer in the other categories is 20.4 on average with a median of 16.9. The distribution of ratings by extremists is clearly more skewed but the difference in the mean range of ratings is less than 1 point on average.

One important difference that emerges from using the model predictions is the much higher relative intensity on guns. In the raw average ratings, the mean range of the ratings for gun issues was 3% higher and the median was 4.5% higher than for other issues. For the modeled ratings the mean is 43% higher and the the median is 67% higher. Respondents are demonstrating through their candidate ratings that they care more about which gun policy is chosen, conditional on the options available, than they report that they do, relative to how much they report they care about which immigration or tax policy is chosen.

The finding that, on average, respondents care more about guns than they let on from their initial ratings masks a more important fact. In general, respondents take fewer factors into account than their unweighted ratings would imply. Figure 6 demonstrates this. The histogram on the left shows the variance in the raw ratings that is accounted for by one issue






**Figure 6:** Concentration of intensity across issue areas. In each panel, the x-axis represents the proportion of the total variance in the policy ratings that is accounted for by each individual’s most variable policy area. The y-axis shows the number of cases. In the left panel, the calculation is based on the raw ratings and in the right panel it is based on the modeled ratings. The red vertical line in each panel is the median.

area. This quantity is bounded from below at  $1/3$  because at that number the ratings vary equally across all three issues. Any number below that implies that the issue area in question does not vary the most. For the median respondent, 48% of the variance is explained by one issue, and most respondents are very close to that number. However the right panel shows the modeled ratings. Once we adjust the ratings for their weight in actual decisions, the variance accounted for by one issue goes up dramatically. Now for the median respondent, 70% of the variance is accounted for by one issue. This “concentration of intensity” is roughly uniformly distributed from 40% to 100%. The modeled ratings show that a single issue plays an outsized role in most respondents’ decisions, but with substantial variance.

Figure 7 visualizes the differences in intensity across individuals. The x-axis represents rightward changes in policy. The leftmost point represents a move from policy 1 to 2, then 2 to 3 and so on. The y-axis represents an absolute change in modeled ratings resulting from a given change in policy. Respondents are divided into deciles according to the mean change in the magnitudes of their ratings, with the most intense (largest magnitude change) decile represented by the darkest line and the least intense decile represented by the lightest line




**Figure 7:** Preference intensity by decile. The x-axis represents incremental rightward changes in policy: from 1 to 2, 2 to 3, up to 6 to 7. The y-axis shows the absolute change in candidate ratings that result. Each line represents a decile or respondents according to the mean absolute change in ratings. The black line shows the average changes in ratings for the decile with the greatest changes (the 10th decile) and the grey line shows the average changes in ratings for the deciles with the smallest changes (the 1st decile) with intermediate shades representing intermediate deciles.

for each policy area. In each area, intensity is highly concentrated. The most intense decile exhibits changes in ratings that range from 2.7 to 4.9 times as big as the 5th decile. This agrees with the finding of the issue publics literature that individuals care about few issues and that the set of people that care about a given issue is relatively small (Converse, 1964; Krosnick, 1990).

Given the survey time required by this measure, it would be fortuitous if respondent intensity could be captured by simpler measures that exist in the literature. One obvious candidate is the classic attitude strength question, asking if the issue in question is important to the respondent personally. Figure 8 breaks out intensity by response to this canonical question, with axes defined as in Figure 7. Unfortunately, responses contain very little information about intensity, with the lines appearing almost indistinguishable.

I should note that it is not obvious how to implement the attitude strength question. In particular, it is not clear how to characterize each issue area. I chose to offer more specificity than a single-word description. Rather than “immigration,” I asked about “treatment of undocumented/illegal immigrants,” a more accurate description, if still incomplete.

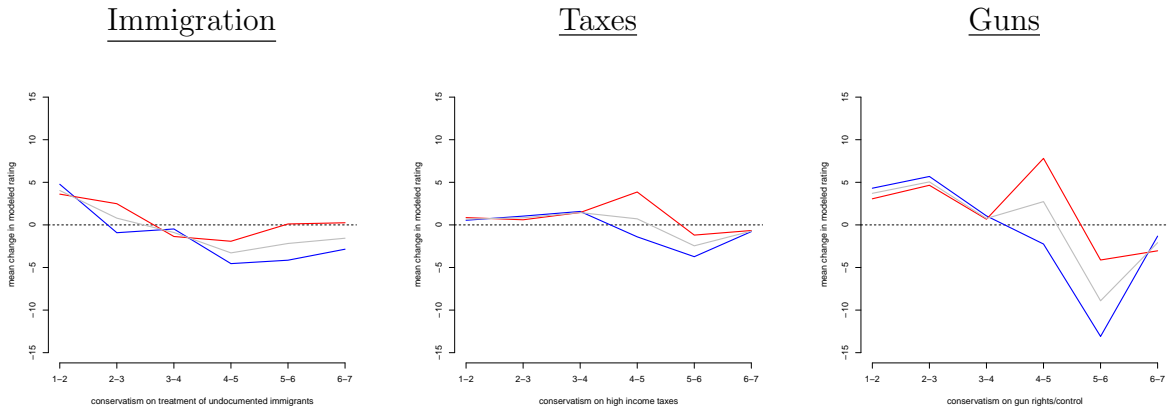


**Figure 8:** Preference intensity by “attitude strength.” The x-axis represents incremental rightward changes in policy: from 1 to 2, 2 to 3, up to 6 to 7. The y-axis shows the absolute change in candidate ratings that result. Each line represents a set of respondents who gave a particular response to the question of whether the policy area in question is important to them personally. The response given by each group is shown in the legend on the top right of each graph.

This raises the possibility that it is my implementation of the question that is causing the measure to have poor explanatory power. For instance, saying that the “treatment of undocumented/illegal immigrants” is important may imply favorable treatment in the minds of respondents. This would explain why the “not at all” responses exhibit greater changes on the right side of the policy spectrum. However, it is not clear what the right way to ask this question is. In its current form, the explanatory power of the attitude strength question is extremely low.

The degree to which changes in candidate positions result in changes in preferences for candidates varies substantially both across and within issue areas. Voters are often close to indifferent over policies that are substantively extremely dissimilar. There are areas of strong agreement and strong disagreement across parties.

To visualize this, Figure 9 shows the *net* change in ratings that result from given changes in candidate positions. This is in contrast to previous graphs that have focused on the average magnitude of the change. The grey line is the overall mean, the red line is the mean for Republicans, and the blue line is the mean for Democrats. The dashed horizontal line



**Figure 9:** Net changes in rating. The x-axis represents incremental rightward changes in policy: from 1 to 2, 2 to 3, up to 6 to 7. The y-axis shows the net change in candidate ratings that result from these changes for each of three groups: Republicans (in red), Democrats (in blue) and the entire sample (in grey).

represents a 0 net change in candidate rating. When the red and blue lines are on the same side of the dashed line, respondents of both parties agree, and when they are on opposite sides they disagree. Points above the line represent support for a conservative change, and points below the line represent opposition to a conservative change.

The red line and the blue line are on the same side of the dotted line at most points, indicating substantial agreement between the two parties. The points at which they are on opposite sides are worth noting. Democrats prefer that taxes on high income people be raised, preferring the lowest available option of 50%. Republicans prefer that taxes on high income people be kept the same. Democrats support an assault weapons ban. Republicans would strongly prefer to “require mandatory training and enhanced background checks for purchase of assault rifles.” For immigration they disagree at point 2-3, but this disagreement is hardly worth noting because the Democrats are so close to the 0 point. Essentially both parties agree that the U.S. should “create a path to citizenship for all undocumented immigrants who have no criminal record” although members of the GOP would prefer that it be limited to those “who are employed, English-speaking... and have been in the US for many years.” Any further moves to the right lose votes on net from both parties.

Just as interesting as these points are the points at which members of either party, or the public at large, are close to indifferent. Conditional on choosing a policy at least as conservative as “Create a legal status but no path to citizenship for undocumented immigrants who were brought to this country as young children,” Republicans are close to indifferent between that option and the two more conservative options, even “Deport all undocumented immigrants immediately.” It would be naive to conclude that Congressional Republicans are currently defying their base by refusing to pass legislation in favor of a path to citizenship. Members of Congress may realize that there is a large degree of indifference on this issue, and that therefore they should use it as a bargaining chip. This can also explain why members with very extreme views on immigration do not get filtered out despite being out of step with the broader electorate.

Similarly, conditional on raising taxes on high income people, members of both parties are close to indifferent about how high those taxes should go- even as high as 100% on the margin. While this might seem outlandish, historically the United States has enacted rates almost this high. These very high tax rates may be difficult for voters to evaluate, putting the issue in the realm of a more Burkean form of representation.

On gun issues, there are large effects throughout the political space, except perhaps in the decision to limit magazine sizes to 10 round or less (3-4). What is remarkable about this issue is that there is a very high degree of intensity among members of both parties, and yet this intensity is coupled with what experts might view as a very narrow range of disagreement. Both parties are strongly against banning guns or banning all but single-round firearms, and both parties are narrowly against limiting clip sizes. This would allow the ownership of the guns used in the overwhelming proportion of cases of gun violence in the United States, including most mass shootings (Lopez, 2018). The United States had an assault weapons ban as recently as 2004. Republicans appear to be strongly in favor of putting strict limits on the ownership of assault weapons.

It is interesting to break respondents out by party because it can help us understand the

dynamic that might emerge from primary elections. If Republicans are indifferent over a set of conservative outcomes and do not vote strategically, than a far-right Republican candidate could easily emerge despite a consensus on the treatment of undocumented immigrants. There are more Democrats in our sample than Republicans. However in every case where the Democrats and Republicans disagree, the overall sample narrowly favors the Republicans. Of course, this need not be the case for issues other than these three.

## 5 Conclusion

In scholarship on representation, Downsian spatial voting has taken on outsize importance. The theory assumes perfect voting and predicts candidate convergence on the median voter. Not surprisingly, it has been found lacking (Ansolabehere, Snyder Jr, and Stewart III, 2001; Fowler and Hall, 2016; Tausanovitch and Warshaw, 2018). Now leading figures in political science can be found advocating for an abandonment of the spatial framework altogether (Achen and Bartels, 2016). However, we often forget that Downs' famous book is for the most part concerned with various extensions of the basic theory (Downs, 1957). We have known for some time that there are reasons not to take our measurements of preferences at face value (Zaller, 1992) and this suggests that perhaps it is our choice of what we measure that is at fault in our inability to find accountability (Gerber et al., 2011).

I propose a fuller approach to measuring preferences that goes beyond ideal points. The approach begins by simply asking respondents to rate a series of policies. Then I model the contribution these ratings make to each individual's ratings of a set of hypothetical candidates. I validate the measure using conjoint analysis.

The measure reveals a number of patterns of immediate interest. Intensity is highly concentrated, as hypothesized by the literature on issue publics (Converse, 1964; Krosnick, 1990). However it is not strongly linked to extremism. There is a gap between how much respondents report that they care about each issue and the role that each plays in their

decisions. There are substantial areas of indifference over substantively different policies. We should not expect strong accountability in these policy regions, a priori. In the political environment at the time of the survey, following closely on the wave of activism that stemmed from the Stoneman Douglas High School shooting, respondents had a much higher level of intensity on average over a set of gun control policies than they did over a set of policies on taxation of high income people and a set of policies regarding the legal status of undocumented immigrants.

These results speak to several existing theories of representation. The directional theory holds that voters prefer more extreme candidates up to some point (Rabinowitz and MacDonald, 1989). This does not appear to be true in general. Respondents often strongly dislike candidates that are distant from themselves, and their choices reflect this. However for some policies there are large regions of indifference and we may observe successful candidates who are fairly extreme on these issues. It is possible that these measures could provide an empirical microfoundation for the theory of thermostatic representation (Wlezien, 1995; Soroka and Wlezien, 2010). The most contested points on each of the three issues under study here are points that are very close to the status quo policy. Intense preferences towards the status quo may have to do with the greater information that voters have about the current status quo and nearby alternatives. This would fit neatly with thermostatic representation. Alternatively, these strong preferences may be a version of the endowment effect (Thaler, 1980). Or they may be endogenous: we have the status quo we have precisely because voters have strong feelings about it.

The endogeneity or exogeneity of preferences is a key question going forward. According to the estimates here, there are a set of policies that appear to be consistently unpopular, and these fit nicely with the reality of which policies are outside the scope of policy debate by elected officials. This suggests a weak form of bottom up democracy. The most preferred policies also fit nicely with those policies espoused by the two major political parties, and the battle lines between the two parties are marked by intense preferences. This could

be because the parties are closely adhering to the preferences of the public. Or it could be because the public comes to care most about what the respective parties support most vociferously (Lenz, 2013). Understanding intensity of preferences could help us understand which views are malleable and which are exogenous “predispositions” (Tesler, 2015).

It is clear that in answering the above questions, measures that conflate intensity and location, such as Likert scales, are inadequate. However this approach could also have important implications for more sophisticated measures of preferences, such as those based on multidimensional scaling (e.g. Clinton, Jackman, and Rivers, 2004). Magnitudes in cardinal scales are almost always determined by assumptions about the shape of utility functions and error distributions that are difficult to test (but see Carroll et al., 2009; Grynaviski and Corrigan, 2006). Theoretically, these magnitudes could be linked to an actual policy outcome (see Bateman, Clinton, and Lapinski, 2017; Oh and Tausanovitch, 2015). However a more straightforward approach would be to use the extent of utility differences as a distance metric. Then the magnitude of distances would have a clear substantive interpretation in terms of relative importance for political choices. This paper provides the means to establish such a metric.

There are many possible applications of the approach presented here, but there are also many possible problems and worthwhile refinements to consider. The model assumes that respondents rate candidates using something akin to a weighted mean of their considerations. Although many steps have been taken to validate the resulting measure, nothing here rules out close alternatives or important individual heterogeneity in decision making. Respondents are only given one shot to rate each policy, and these ratings are the basis for all that follows. There are likely many avenues for doing better. Each policy area was rated on only seven discrete points. There may be many respondents for which large changes in preference occur over points that were “skipped over.” Preference functions are continuous but the policies that were rated here are discrete. These are only a few examples of areas for improvement that could be addressed in future work.



Establishing the external validity of this approach is another important task for future work. All of the evidence so far comes from experiments with hypothetical choices, and only three policies have been studied. One obstacle is that voters increasingly evaluate the policy platforms of parties, not individual candidates (e.g. Bonica and Cox, 2018). This reduces the variation available to test the effects of preference intensity in the context of real U.S. elections. It may be necessary to find a more amenable context for testing this approach against observational data.

## References

- Achen, Christopher H, and Larry M Bartels. 2016. Democracy for realists: Why elections do not produce responsive government. Princeton University Press.
- Adams, James, and Samuel Merrill III. 2003. “Voter turnout and candidate strategies in American elections.” The Journal of Politics 65(1): 161–189.
- Ansolabehere, Stephen, James M Snyder Jr, and Charles Stewart III. 2001. “Candidate positioning in US House elections.” American Journal of Political Science pp. 136–159.
- Arrow, Kenneth J. 1951. Social Choice and Individual Values. New York: John Wiley & Sons.
- Ballard-Rosa, Cameron, Lucy Martin, and Kenneth Scheve. 2017. “The structure of American income tax policy preferences.” The Journal of Politics 79(1): 1–16.
- Bateman, David A, Joshua D Clinton, and John S Lapinski. 2017. “A House Divided? Roll Calls, Polarization, and Policy Differences in the US House, 1877–2011.” American Journal of Political Science 61(3): 698–714.
- Bechtel, Michael M, and Kenneth F Scheve. 2013. “Mass support for global climate agreements depends on institutional design.” Proceedings of the National Academy of Sciences 110(34): 13763–13768.
- Bishin, Benjamin. 2009. Tyranny of the minority: The subconstituency politics theory of representation. Temple University Press.
- Bonica, Adam, and Gary W Cox. 2018. “Ideological Extremists in the US Congress: Out of Step but still in office.” Quarterly Journal of Political Science Forthcoming.
- Carroll, Royce, Jeffrey B Lewis, James Lo, Keith T Poole, and Howard Rosenthal. 2009. “Comparing NOMINATE and IDEAL: Points of difference and Monte Carlo tests.” Legislative Studies Quarterly 34(4): 555–591.

- Clinton, Joshua, Simon Jackman, and Douglas Rivers. 2004. "The statistical analysis of roll call data." American Political Science Review 98(2): 355–370.
- Converse, Philip E. 1964. "The Nature of Belief Systems in Mass Publics." In Ideology and Discontent, ed. David Apter. New York: Free Press.
- Coppock, Alexander, and Oliver A. McClellan. 2017. "Validating the Demographic, Political, Psychological, and Experimental Results Obtained from a New Source of Online Survey Respondents." Unpublished manuscript .
- Downs, Anthony. 1957. An Economic Theory of Democracy. Harper.
- Fiorina, Morris P. 1974. Representatives, roll calls, and constituencies. Lexington Books.
- Fowler, Anthony, and Andrew B Hall. 2016. "The Elusive Quest for Convergence." Quarterly Journal of Political Science 11(1): 131–149.
- Gerber, Alan S, Gregory A Huber, David Doherty, and Conor M Dowling. 2011. "Citizens' policy confidence and electoral punishment: A neglected dimension of electoral accountability." The Journal of Politics 73(4): 1206–1224.
- Grynaviski, Jeffrey D, and Bryce E Corrigan. 2006. "Specification issues in proximity models of candidate evaluation (with issue importance)." Political Analysis 14(4): 393–420.
- Hainmueller, Jens, Daniel J Hopkins, and Teppei Yamamoto. 2013. "Causal inference in conjoint analysis: Understanding multidimensional choices via stated preference experiments." Political Analysis 22(1): 1–30.
- Kahneman, Daniel, Peter P Wakker, and Rakesh Sarin. 1997. "Back to Bentham? Explorations of experienced utility." The quarterly journal of economics 112(2): 375–406.
- Krosnick, Jon A. 1990. "Government policy and citizen passion: A study of issue publics in contemporary America." Political behavior 12(1): 59–92.

- Krosnick, Jon A, and Richard E Petty. 1995. "Attitude strength: An overview." Attitude strength: Antecedents and consequences 1: 1–24.
- Lenz, Gabriel S. 2013. Follow the leader?: how voters respond to politicians' policies and performance. University of Chicago Press.
- Lopez, German. 2018. "Why an assault weapons ban can't address America's gun problem." Vox.com (Apr).
- Oh, Jason S, and Chris Tausanovitch. 2015. "Quantifying Legislative Uncertainty: A Case Study in Tax Policy." Tax L. Rev. 69: 485.
- Rabinowitz, George, and Stuart Elaine Macdonald. 1989. "A directional theory of issue voting." American Political Science Review 83(1): 93–121.
- Riker, William H, and Peter C Ordeshook. 1968. "A Theory of the Calculus of Voting." American political science review 62(1): 25–42.
- Soroka, Stuart .N., and Christopher Wlezien. 2010. Degrees of Democracy: Politics, Public Opinion, and Policy. Cambridge University Press.
- Tausanovitch, Chris, and Christopher Warshaw. 2018. "Does the Ideological Proximity Between Candidates and Voters Affect Voting in U.S. House Elections?" Political Behavior 40(1): 223–245.
- Tesler, Michael. 2015. "Priming predispositions and changing policy positions: An account of when mass opinion is primed or changed." American Journal of Political Science 59(4): 806–824.
- Thaler, Richard. 1980. "Toward a positive theory of consumer choice." Journal of Economic Behavior & Organization 1(1): 39–60.
- Tomz, Michael, Jessica LP Weeks, and Keren Yarhi-Milo. 2018. "Public Opinion and Decisions about Military Force in Democracies." Manuscript .

Wlezien, Christopher. 1995. "The public as thermostat: Dynamics of preferences for spending." American journal of political science pp. 981–1000.

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

## A Data representativeness

Table 2 compares the demographic composition of the survey used for this study with the demographic composition of the face-to-face and online versions of the 2016 American National Election Study, the gold standard survey in political science. The composition is quite similar.

Table 3 shows the answer to a question about whether the United States is headed “in the right direction” by gender and race, taken from the survey used for this study. Table 4 shows the same information, but taken from a contemporaneous Economist/YouGov survey. Tables 5 and 6 compare answers by party and race. The results are again similar.

	This survey	ANES 2016 FTF	ANES 2016 WEB
Female	53%	53%	53%
18-29	19%	17%	15%
30-44	28%	24%	26%
45-64	34%	36%	37%
65+	20%	25%	22%
White	74%	68%	73%
Black	10%	10%	9%
Hispanic	6%	14%	9%
Other	10%	8%	8.4%
Dem	37%	31%	35%
Ind	34%	42%	35%
Rep	29%	27%	30%
Did not graduate high school	3%	8%	6%
High school graduate	20%	22%	18%
Some college	25%	17%	23%
2-year college degree	12%	16%	14%
4-year college degree	28%	23%	23%
Postgraduate degree	11%	13%	17%

**Table 2:** Demographic comparison between the survey used for this study and the face-to-face and online versions of the 2016 American National Election Study.

	Male	Female	18-29	30-44	45-64	65+
Generally headed in the right direction	39%	30%	25%	36%	36%	37%
Off on the wrong track	56%	58%	63%	54%	56%	56%
Not sure	5%	12%	13%	10%	8%	6%
N	660	745	262	392	480	279

**Table 3:** Response to question “Would you say things in this country today are...” on the survey used for this paper, April 4 to 11.

	Male	Female	18-29	30-44	45-64	65+
Generally headed in the right direction	34%	30%	24%	23%	37%	42%
Off on the wrong track	54%	54%	54%	60%	52%	50%
Not sure	13%	16%	22%	17%	11%	9%
N	647	853	157	390	643	310

**Table 4:** Response to question “Would you say things in this country today are...” on the YouGov Economist Poll, April 8 to 10.

	White	Black	Hispanic	Other	Dem	Ind	Rep
Generally headed in the right direction	37%	17%	27%	33%	13%	27%	68%
Off on the wrong track	54%	73%	64%	58%	81%	58%	25%
Not sure	9%	10%	9%	9%	6%	14%	7%
N	1,051	140	85	138	517	483	414

**Table 5:** Response to question “Would you say things in this country today are...” on the survey used for this paper, April 4 to 11.

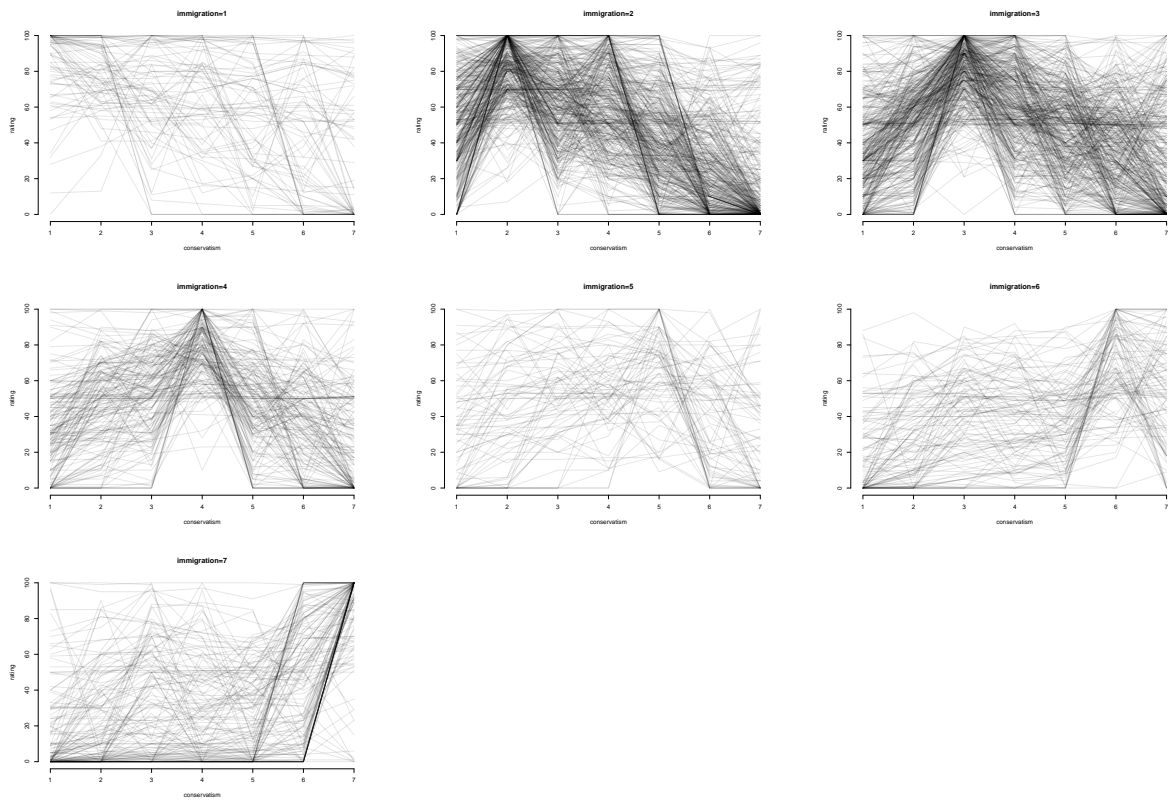
	White	Black	Hispanic	Other	Dem	Ind	Rep
Generally headed in the right direction	37%	16%	26%	11%	13%	24%	72%
Off on the wrong track	51%	64%	58%	56%	76%	55%	20%
Not sure	12%	21%	16%	33%	11%	20%	8%
N	1,109	173	146	72	548	593	359

**Table 6:** Response to question “Would you say things in this country today are...” on the YouGov Economist Poll, April 8 to 10.

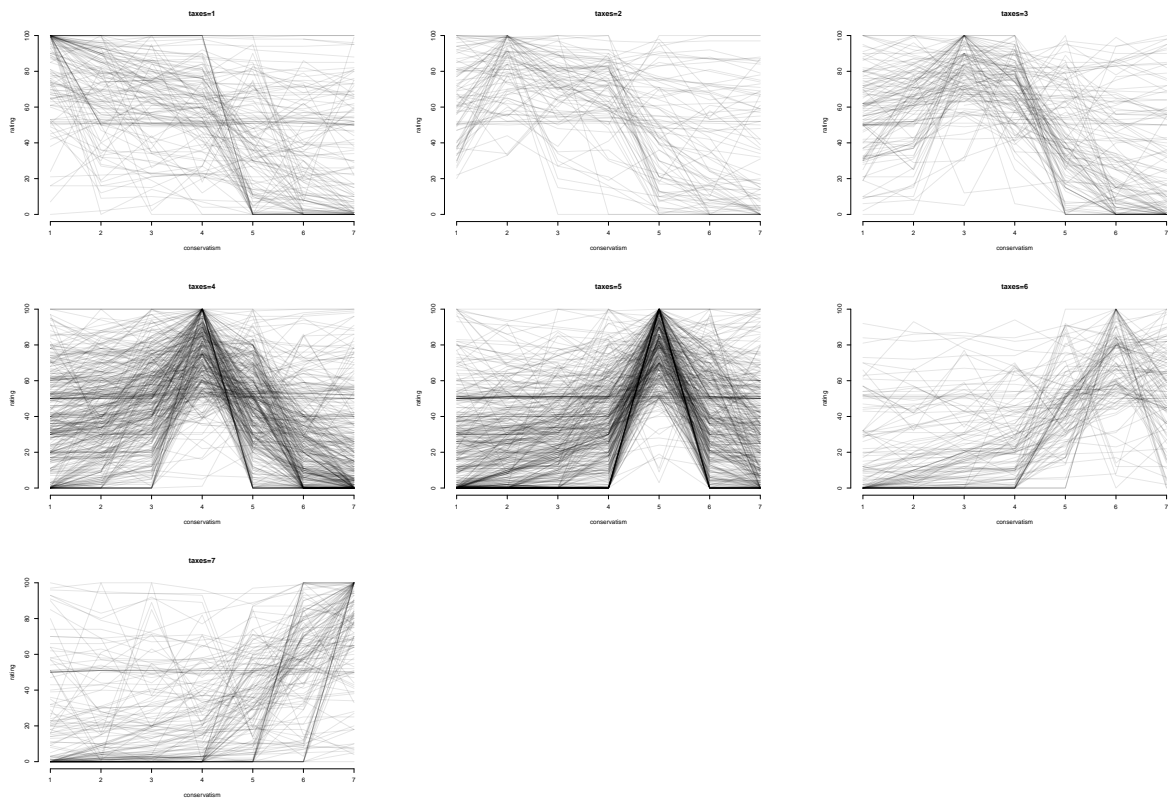
## **B Individual self-reported ratings**

Figures 10, 11 and 12 show each individual self reported rating by most preferred policy for each issue area.

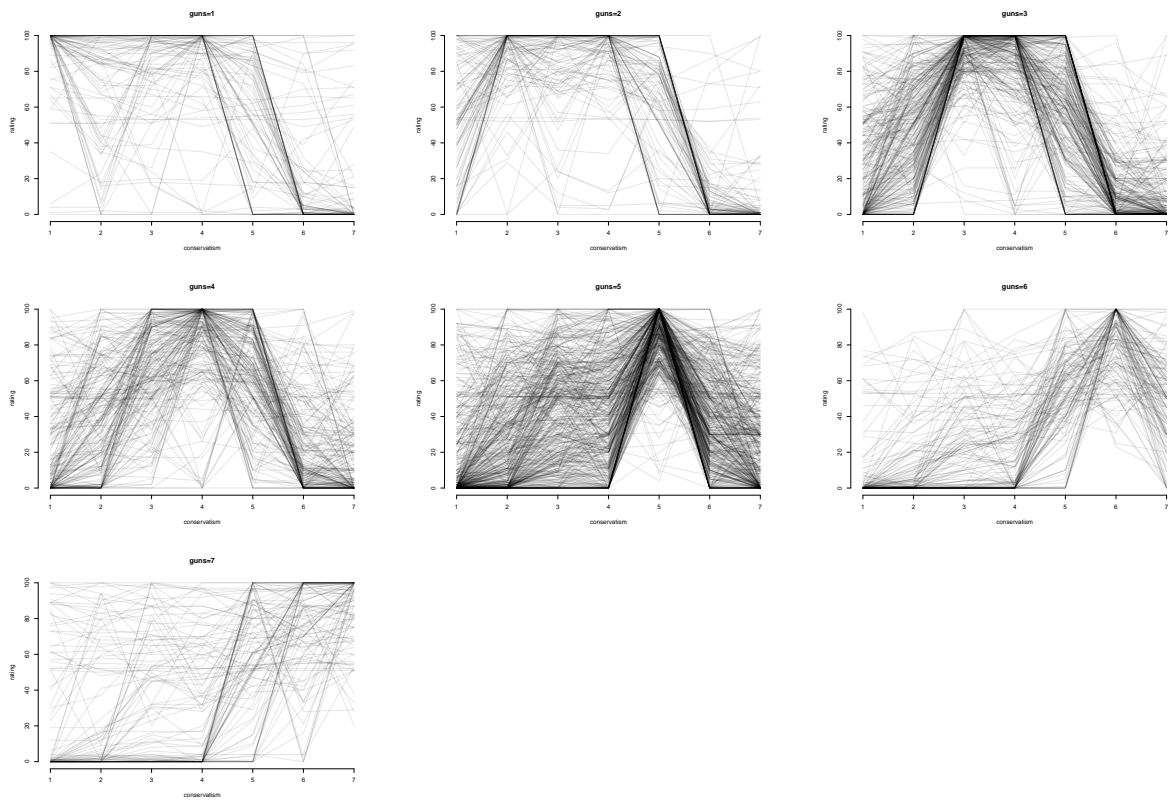




**Figure 10:** Individual self-reported ratings on immigration, by most preferred policy.




**Figure 11:** Individual self-reported ratings on taxes, by most preferred policy.



**Figure 12:** Individual self-reported ratings on guns, by most preferred policy.

# C Model results

Figure 13 shows the predicted effects of policy positions on candidate ratings for each possible most preferred policy in each issue area. This figure uses the main specification from the paper to generate predictions.

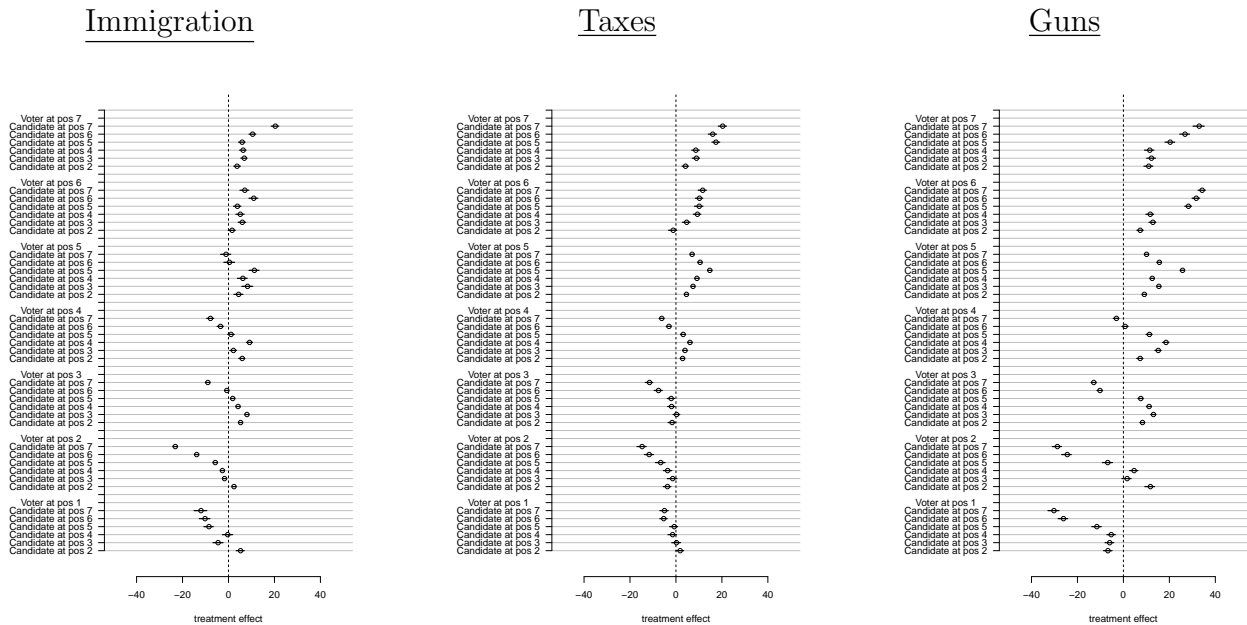


**Figure 13:** Predicted effects of policy positions on ratings, by most preferred policy. For each policy area this graph shows the predicted effects on ratings of candidate positions, relative to a baseline of the most liberal position (1). These effects are grouped by respondents according to their most preferred position.


## D Conjoint results

Figure 14 shows the predicted effects of policy positions on candidate ratings for each possible most preferred policy in each issue area. These effects are calculated using conjoint analysis.

Figure 15 shows the effect of policy positions on the probability of choosing a candidate, also calculated using conjoint analysis.



**Figure 14:** Treatment effects of policy positions on ratings, by most preferred policy. For each policy area this graph shows the effects of candidate positions on ratings, relative to a baseline of the most liberal position (1), estimated by conjoint analysis. These effects are grouped by respondents according to their most preferred position. Bars represent 95% confidence intervals.



**Figure 15:** Treatment effects of policy positions on choice probabilities, by most preferred policy. For each policy area this graph shows the effects of candidate positions on choice probabilities, relative to a baseline of the most liberal position (1), estimated by conjoint analysis. These effects are grouped by respondents according to their most preferred position. Bars represent 95% confidence intervals.

## E Alternative models

The model used in this paper is a straightforward linear regression, represented by equation 1. However simpler and more complex alternatives are available. One option would be to dispense with the separate coefficients for each issue area. In this model, each respondent has an individual-specific slope and intercept, but the effects of the ratings for not vary across issues or candidate party. I call this the “pooled ratings model”:

$$rating_i = \gamma_0 + \gamma_1(imm\_rating_{imm_i} + tax\_rating_{tax_i} + gun\_rating_{gun_i} + party\_rating_{party_i}) + \epsilon_i$$


A more complex model could borrow power across respondents by using random coefficients. The random effects model is the same as the original model but with the intercept

and coefficients drawn from a common normal distribution with a variance that is estimated from the data. If people are indexed by  $j$  then this can be written as:


$$\begin{aligned}
 rating_{ij} &= \beta_{0j} + \beta_{1j}imm\_rating_{imm_{ij}} + \beta_{2j}tax\_rating_{tax_{ij}} \\
 &\quad + \beta_{3j}gun\_rating_{gun_{ij}} + \beta_{4j}party\_rating_{party_{ij}} + \epsilon_{ij} \\
 \beta_{0j} &\sim N(\beta_0, \sigma_0) \\
 \beta_{1j} &\sim N(\beta_1, \sigma_1) \\
 \beta_{2j} &\sim N(\beta_2, \sigma_2) \\
 \beta_{3j} &\sim N(\beta_3, \sigma_3) \\
 \beta_{4j} &\sim N(\beta_4, \sigma_4)
 \end{aligned}$$

Figure 16 compares these alternative models to the estimated conjoint effects. The pooled ratings model substantially underperforms the specification in the paper, explaining only 72% of the variation in the conjoint effects as opposed to 81% for my preferred model. The estimated effects are also smaller, at 57% of the conjoint effects as opposed to 67%. The random effects model is more complex and is also marginally worse, with an  $R^2$  of 80% and effects that are 63% as big as the conjoint effects.

Pooled Ratings Model



Random Effects Model



**Figure 16:** Comparison of modeled ratings to conjoint results. Each point corresponds to the set of respondents who chose a particular most preferred policy in a given issue area and the estimated effect of a given candidate policy position on their candidate rating. These effects are calculated relative to a baseline where the candidate takes position 1, the most liberal position. The y-axis is the effect calculated using conjoint analysis and the x-axis is the effect calculated using the average predicted values from a model, where the independent variables are the respondents' ratings of candidate policies and party. The model used in the left panel does not allow the coefficients to differ by the party or policy rated. In the right panel the model allows the coefficient to differ by policy area and for the party rating and uses random effects. The dashed line is the 45-degree line and the black line is the regression line. Blue dots are effects on gun policy, red dots are effects on immigration, and green dots are effects on taxes.




## F Candidate choice and abstention

The classic finding in the literature on representation does not depend on intensity at all. If voters have single-peaked preferences and choose the candidate they prefer more, then the position of the median voter is a dominant choice for candidates. In such a world it is sufficient to know each respondents' most preferred policy and not consider intensity.

In this section I show that candidate choices are stochastically related to the difference in candidate ratings. The setting is admittedly idealized, but it shows that at least in a hypothetical case, candidate choices are probabilistic and these probabilities are effected by the voter's intensity with respect to the candidates. I also show that abstention is related to intensity in a manner that is consistent with the literature. Under these circumstances, candidates have an incentive to cater to the most intense voters.

Recall that voters are asked to rate 20 candidates, and then these candidates are paired up at random and voters are asked to choose between them. Each candidate is included once. As we might expect, these choices are made with substantial error, and the degree of difference in ratings is highly predictive of the error rate. Figure 17 shows the logistic regression curve from a model of candidate choice based on the difference in ratings for all candidates, pooled over all respondents. As the candidate in question (chosen arbitrarily) becomes more and more advantaged in terms of her rating, the probability that she is chosen becomes greater. Not too surprisingly, the probability of choosing a given candidate is almost exactly 50% when the two candidates are equally rated, and the change in probabilities is almost perfectly symmetric given that the choice of which candidate is 0 and which is 1 is arbitrary. When a candidate is rated 25 points higher than the other candidate, she has a 76% chance of being chosen, leaving a 24% chance of choosing the lower rated candidate.

The existence of some error is crucial, because it means that the median is no longer automatically the winning point. This analysis also lends credence to the measure of intensity by showing that it predicts choices between candidates. Voting probabilities among more intense respondents are much more responsive to changes in policy positions. For this reason,



**Figure 17:** Candidate choices as a function of the difference in ratings. The predicted probability of choosing an arbitrarily designated candidate is on the y-axis. The difference in ratings between the two candidates is on the x-axis. The vertical lines represent the distribution of differences in ratings for the candidates who are not chosen (on the bottom) and the candidates who are chosen (on the top). Error bounds are represented by the dotted lines.

legislators have an incentive to give greater weight to the preferences of the more intense respondents

Of course, this analysis has serious limitations, and should not be taken literally. External validity is an obvious concern. We expect much higher error rates in choices made in seconds between abstract candidates than we do in votes cast for real candidates who conduct extensive media campaigns over many months. This analysis demonstrates the mere existence of error and that this error is related to the difference in ratings.

The other major limitation is that unlike the model of ratings, we do not have enough statistical power to model choices at the individual level. This is a serious drawback. There is reason to believe that voters differ in their *acuity*. Some will perceive very small differences in ratings and act accordingly, while others will only perceive big differences. This alters the relationship between candidate positions and vote probabilities. For applied research it may be necessary to have an understanding both of intensity and acuity to have firm predictions


about the effects of candidate positions or attributes.

Without answering this question definitively, one way to assess the scope of the difference that acuity might make is to compare the modeled effects on *ratings* to a conjoint analysis of the candidate *choices*. This analysis captures both differences in intensity and in acuity, although it can only be conducted at the group level. This will mask compositional effects but give us a sense of whether the modeled ratings capture overall changes in probability of selection.

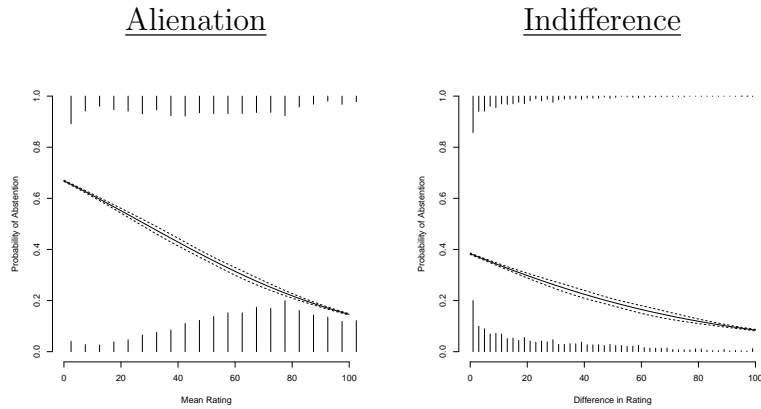
Figure 18 shows the results using the same approach that was used for the conjoint analysis of ratings. The resulting model fit is surprisingly strong. A 1 point change in the modeled rating corresponds with a 2 percentage point change in support for a given candidate. The modeled ratings explain 82% of the effect on choices. The strength of this effect suggests, contrary to expectations, that aggregate predicted changes in ratings correspond closely to changes in the probabilities of selecting a candidate without accounting for differences in respondent acuity. It is possible that acuity and intensity are strongly related, that differences in acuity cancel each other out within groups, or that acuity is more homogeneous than anticipated.

One of the more important topics that these measures could address is the question of whether candidates are more incentivized to mobilize their base or to win over swing voters. A key part of this puzzle is understanding the link between intensity and turnout. As a baby step towards this goal, half of the candidate choice questions include a “neither” option, allowing the respondent to effectively abstain. There are two well-known sources of abstention based on policy differences with candidates. Voters may abstain because of indifference between the two candidates if they do not have a strong preference for one over the other. Or they may abstain if they dislike both candidates sufficiently. Even if the better candidate is strongly preferred, voters may get disutility from voting for a candidate that they dislike. This has been called abstention due to alienation (Adams and Merrill III, 2003).

I conduct a simple logit analysis, modeling the choice to abstain as a function of the



**Figure 18:** Comparison of modeled ratings to conjoint results. Each point corresponds to the set of respondents who chose a particular most preferred policy in a given issue area and the estimated effect of a given candidate policy position on their either their candidate choice (y-axis) or candidate rating (x-axis). These effects are calculated relative to a baseline where the candidate takes position 1, the most liberal position. The y-axis is the effect calculated using conjoint analysis and the x-axis is the effect calculated using the average predicted values from a set of linear regressions for each respondent, where the independent variables are the respondents' ratings of candidate policies and party. The black line is the regression line. Blue dots are effects on gun policy, red dots are effects on immigration, and green dots are effects on taxes.



**Figure 19:** Abstention due to alienation and indifference. The left panel shows predicted values for abstention in a choice between two candidates as a function of the mean rating of the candidates. The right panel shows predicted values for abstention as a function of the difference in ratings between the two candidates. The vertical lines show the distribution of the x-axis values for the abstainers, at the top of the graph, and the non-abstainers, at the bottom.

difference in ratings between the two candidates and the mean rating of the two candidates. The results for the two independent variables are graphed in Figure 19. The left panel of the graph shows that when both candidates are close to a 0 rating, the probability of abstention is almost 70%. When the candidates approach a mean rating of 100, the probability of abstention declines to 15%. These numbers underestimate the resulting gap in total abstention because they hold the difference between candidates at the mean, an assumption that is clearly implausible as the mean rating of the two candidates approaches the 0 lower bound and 100 upper bound. The right panel shows that when there is 0 difference between the candidates, the predicted probability of abstention is 38%. The abstention probability bottoms out at 8% as the difference in ratings approaches 100. These results are clear confirmations of abstention due to alienation and indifference, albeit in a highly idealized setting. These patterns lend further credence to the underlying measure.

The picture of decision making that is painted by these measures is very different from the standard model of simple proximity and perfect voting. The degree to which changes in candidate positions result in changes in “utility” for candidates varies substantially both

across and within issue areas.