The institutional sources of policy bias:*

Gabor Simonovits†

August 7, 2017

JOB MARKET PAPER

Abstract

How well does public policy represent mass preferences? The approaches typically employed in empirical research on this important question fall short for two reasons. First, they fail to measure how much more liberal policies are compared to preferences. Second, they do not assess the heterogeneity of preferences within jurisdictions, and thus do not consider how the quality of representation depends on the level to which policy decisions are delegated to. Here we overcome both of these problems by generating estimates of Americans' preferences on the minimum wage which are measured on a scale that is comparable to observed policies and describe low levels of geographic aggregation. Using these estimates, we demonstrate that most people are poorly represented by state minimum wage laws because of two key reasons. First, in each state, the minimum wage is much lower than the average rate preferred by state residents, leading to a pronounced bias against the preferences of the poor. Second, because preferences vary within states to a great deal, they are difficult to match by a single policy even in the absence of an overall policy bias. While minimum wage laws in the U.S. are typically set by elected officials and cover entire states, our results show that policies brought about by direct democratic institutions and at more local levels reflect preferences substantially better. These findings suggest that standard data and measures yield incomplete evidence about the relationship between public opinion and policy in the U.S..

*This paper is partly based on joint work with Andrew Guess and Jonathan Nagler. I thank my co-authors' permission to use these results here. I also thank Neal Beck, Jim Bisbee, Pat Egan, Sean Kates, Neil Malhotra, Arthur Spirling and Stephanie Zonszein for helpful suggestions. The survey data presented in the paper were collected by the Social Media and Political Participation Lab at New York University, and is used with the permission of its PI, Jonathan Nagler. All other data used in this paper is freely available online.

†Ph.D. Candidate, New York University, Department of Politics, 19 West 4th St, NYC, NY, 10012, email: simonovits@nyu.edu

How well does public policy represent mass preferences in U.S states? Political scientists approach this question with strong theoretical expectations. Classic accounts of political competition predict that electoral incentives force lawmakers to enact centrist policies that are close to the preferences of their states' median or average voters (Downs, 1954; see also Erikson, 2014). In the context of federalist systems, this means that the variation of policy outcomes across subnational units is expected to reflect the geographically heterogeneous preferences of voters. The U.S. federalist system is thus presumed to lead to welfare gains because like-minded voters are sorted into subnational constituencies that enact policies close to average preferences within these constituencies (Tiebout, 1956; Oates, 1999).

However, evaluating whether these expectations hold turns out to be surprisingly difficult using standard empirical approaches to the study of representation. The findings from this work can largely be summarized by two stylized facts. First, many studies have demonstrated that policies across the U.S. are *responsive* to public opinion in the sense that more liberal states tend to adopt more liberal policies (e.g. Erikson et al., 1993). But second, a newer line of evidence suggests that state policies are nevertheless often *incongruent* with public opinion in that they fail to correspond with majority preferences (e.g. Matsusaka, 2010; Lax and Phillips, 2012).

Valuable as these findings are, studies employing either responsiveness or congruence as evaluative criteria cannot test some central claims about the quality of representation in the U.S. federal system. First, because they do not compare preferences and individual policies on a scale that cor-

responds to the standard spatial account of politics, they fail to provide substantively meaningful estimates of policy *bias*: the distance between opinion and policy within states. Second, because they typically ignore the heterogeneity of preferences within states, they cannot assess how well policies represent the entire distribution of preferences, rather than just that of the average voter. Importantly, a lack of attention to how much preferences vary within states—especially as compared to within larger and smaller jurisdictions—precludes an evaluation of whether the delegation of policy-making powers to sub-national governments under federalism actually improves representation.

We discuss these shortcomings in more detail later, but here we provide a brief summary of our argument using hypothetical configurations of preferences and policies. Figure 1 provides intuition as to why both responsiveness and congruence can mismeasure the extent to which public policy corresponds with preferences. In the example depicted in Panel A, policy is highly responsive to average preferences across states (the correlation is .98). But policies are nevertheless very far from preferences within states (with policy bias equivalent to two-thirds of the range of average preferences). In Panel B, policy in a hypothetical state is congruent with majority opinion compared to the status quo. But this policy too is biased, as it is in fact located far from the ideal point of the state's median voter. Using either measure, bias—the deviation between policy and average preferences—remains unobserved to the researcher.

Figure 1: Inferential challenges to the study of representation

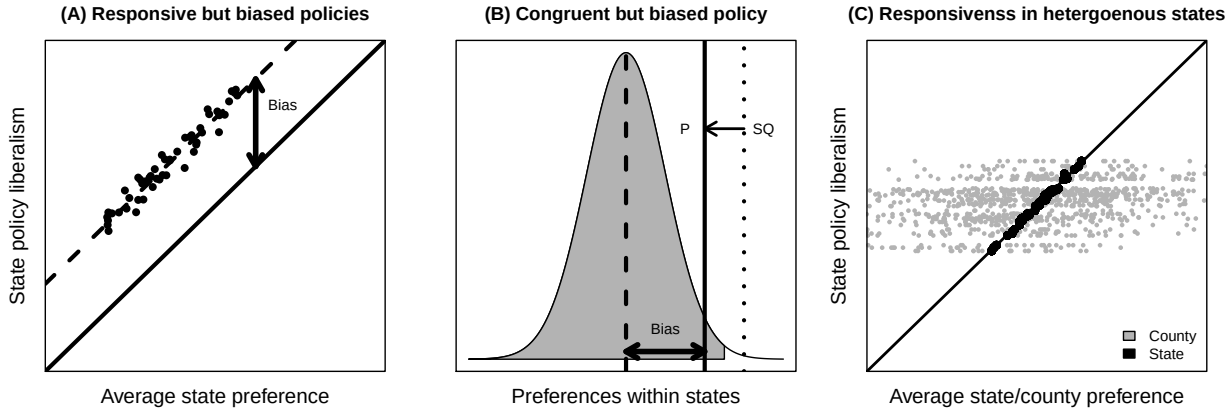

Note: **Panel A** depicts the hypothetical relationship between average preferences and policies across a set of states, measured on different scales. The dashed indicates the best fitting line (observed to the researcher) and the solid line depicts the unbiased mapping (i.e. the 45-degree line) between the two variables which remains unobserved to the researcher. Policy is highly responsive to public opinion but deeply biased. **Panel B** depicts the distribution of ideal points in a hypothetical state as well as the the ideological position of two alternative policies: an enacted policy P (solid line) and the status quo (dotted line). A large share of the electorate (shaded grey) prefers P to the status quo, and thus it is congruent with majority opinion. But the policy is deeply biased: it is far from the preferences of the median voter (dashed line). **Panel C** depicts the hypothetical relationship between enacted state policies and average preferences in counties across a set of states. Grey dots denote average preferences in counties. Counties are nested in states, whose average preferences are denoted with black dots. The solid line depicts the unbiased mapping (i.e. the 45-degree line) between opinion and policy. Even in a federal system where policy perfectly reflects average preferences across states, state policies can be deeply biased with respect to the preferences of many state residents.

Panel C of the figure illustrates the consequences of ignoring preference heterogeneity across different subnational levels. It depicts a stylized example in which states (black dots) are relatively similar in terms of average preferences, but average preferences of county residents (grey dots) vary greatly within states. In this scenario, a focus limited to average preferences across states will conclude that the system is perfectly responsive and unbiased to public opinion. But it will fail to uncover the fact that state policies are in fact located quite far from most voters' preferences. This example also shows that assessments of responsiveness can vary widely

depending on the level of analysis: a strong relationship between policy outcomes and average state preferences can mask a nearly non-existent one between the same policies and preferences measured at lower levels of geography.

The goal of this paper is to introduce new measures, data and methods that can overcome these challenges and to demonstrate the value of these approaches for empirical studies of the opinion-policy relationship. On the one hand, we draw on the theoretical approach introduced by Achen (1978) and conceptualize the quality of representation based on the ideological distance between individual preferences and policies. On the other hand, we build on recent advances in statistical modelling (Lax and Phillips, 2009; Ghitza and Gelman, 2013) in order to estimate of local preferences at extremely low levels of aggregation.

Our focus is the relationship between opinion and policy across states on the minimum wage—the hourly rate employers are required to pay to workers by law. The minimum wage is a particularly rich case for the study of subnational representation. It is a highly salient, easy-to-understand policy about which it is reasonable to expect individuals to form valid opinions. We can quantify both preferences and policies on the same scale (wages in dollars per hour), permitting substantively meaningful comparisons between them. Minimum wage laws are enacted at the federal, state and local levels, so identifying the level of jurisdiction at which these laws maximize welfare is highly relevant for the study of representation in the U.S.

Our empirical analysis relies on two novel empirical approaches to explore the representation of constituent preferences. First, we fielded a nationally representative survey in which we measured individual preferences about the minimum wage using an open-ended question. The resulting data was used to generate estimates of average preferences in each state that are on the same scale as policy outcomes both across and within states. Second, we used data on votes cast on minimum-wage referenda in four states to estimate average preferences at the county and city level with a structural measurement model. These extremely localized preference estimates allowed us to explore the consequences of preference heterogeneity within states for representation when laws are made at the state level.

We report two key findings. First, using our data on state-level preferences and policy outcomes, we show that even though there is an almost perfect correlation between preferences and policy across states, such responsiveness coexists with a remarkably large conservative policy bias. On average, state minimum wages are set at a level approximately two dollars per hour lower than the wage state residents would prefer, exceeding the difference between the average preferred rate in the nation's most liberal state (New York, $11.70 per hour) and the most conservative state (Oklahoma, $9.44 per hour). As a result of this conservative policy bias, state minimum wages are lower than those preferred by all but the most affluent Americans, providing another instance of policy bias in favor of the wealthy (Gilens 2014; Gilens and Page, 2015).

Second, using our estimates of county-level preferences we demonstrate that the variation of preferences within states often dwarfs the differences

across states. As a result, even if the adopted minimum wage laws were to reflect average preferences within states, they would still be distant from many state residents' preferences. Minimum wage laws are particularly distant from the preferences of residents of atypical counties, such as conservative rural areas in rich states (which oppose virtually any increase in the minimum wage) or liberal urban centers in poor states (which prefer dramatic minimum wage hikes).

Our findings also point to two promising ways to improve the quality of representation, at least on this issue. First, we found that policy outcomes are substantially closer to average preferences in states with access to direct democracy and successful initiatives that increased the minimum wage reduced bias by more than 50% on average. Second, our results also indicate that delegating policy decisions to the level of cities rather than states leads to polices that better reflect mass opinion in the local level. Taken together these results suggest that institutional changes can successfully reduce the gap between mass preferences and policies: direct democracy can re-center policy outcomes that have drifted towards the preferences of the affluent, while decentralization can guarantee the representation of cities demanding more liberal policies compared to the states in which they are located.

All told, our findings make a strong case for the adoption of these kinds of measures, methods and data by empirical scholars of representation as they explore the opinion-policy relationship. Current methods may be reaching inappropriately sanguine conclusions about the quality of representation in the U.S. and the proper level of government at which laws maximize

voter welfare. Expanding the empirical toolkit along the lines illustrated here can yield stronger evaluations of the key arguments for American federalism.

Conceptual overview

Empirical studies of representation are grounded in a populist account of democracy (e.g. Dahl, 1971; Achen and Bartels, 2016) that evaluates policy outcomes based on how closely they reflect constituent preferences (Achen, 1978). The existing literature has been confronted with two challenges when making empirical inferences about the opinion-policy relationship. First, the lack of data on policy and preferences that are measured on the same scale has precluded direct and meaningful assessment of the deviation of policy from opinion. Second, the focus of empirical research on the central tendency of within-jurisdiction preferences – but *not* their variation – has made it impossible to assess how much the decentralization of policy decisions improves representation vis-a-vis more centralized policy-making. In this section we first consider these two limitations and then introduce our proposed solution to these problems.

Responsiveness and congruence as measures of representation

Researchers have responded to the lack of data on policy and preferences measured on the same scale with two distinct approaches. On one hand, studies exemplified by Erikson and his co-authors (1993) have explored policy *responsiveness*: the relationship between public opinion and policy outcomes *across* states. The key idea behind this approach is that even if preferences and policy outcomes are measured on different scales, responsiveness is informative about whether more liberal states end up

with more liberal policies. Studies of responsiveness using a variety of approaches to measure both policies and preferences have found exactly this pattern (Erikson et al., 1993; Lupia et al., 2010; Gerber, 1996; Norrander, 2000; Caughey and Warshaw, 2016).

On the other hand, studies focusing on individual policy outcomes (Matsusaka, 2010; Lax and Phillips, 2009; 2012) have explored *congruence*: whether or not individual policies in a jurisdiction are supported by the majority. These studies have led to a much less optimistic conclusion about statehouse democracy: they find that states often fail to enact policies supported by even large opinion majorities, and congruence is barely greater than would be expected by chance (Lax and Phillips, 2012, p. 149). Thus, a key issue in the study of the opinion-policy relationship is an apparent tension between the results of these two approaches (Erikson, 2015).

One way to reconcile the seemingly contradictory findings is to realize that both approaches test necessary rather than sufficient conditions of the closeness of policy to average opinion. First, to the extent that average preferences vary across states, in order for policies to reflect these differences they need to "respond" to mass opinion (a necessary condition). However, responsiveness is not a sufficient condition for policies to be close to average preferences: as we illustrate in Figure 1a, responsiveness can coexist with bias, resulting in policies that are far from mass preferences.[1]

The middle panel of the figure illustrates a similar inferential challenge arising from the use of congruence. The figure shows the distribution of

[1]The idea that studies that measure the general liberalism of mass preferences and policies on different scales cannot assess the "distance" of policies and preferences within states has been pointed out by many (e.g. Achen, 1978; Matsusaka, 2001; Lax and Phillips, 2012). At the same time, the use of responsiveness to make normative claims about representation is still in use even if these limitations are noted.

preferences within a state and depicts two policy alternatives: the status quo (dashed line) and a proposed policy (solid line). Because the solid line is closer to the ideal point of the average voter, the proposal would garner majority support in the state. In fact, under standard assumptions we would expect that all voters to the left of the cut-point between the two policies would prefer the proposal to the status quo. Note, however, that the proposed policy is still extremely conservative compared to the preferences of the state.[2] Thus, it would be misleading to conclude that enacting the proposal over the status quo would lead to a policy that is close (as opposed to just closer) to average preferences. To summarize, neither the association of preferences and policies across states, nor the congruence of a policy with the majority opinion implies that policies are close to average preferences.

The crucial role of preference heterogeneity within states

A distinct issue with the current approaches to the study of representation is their exclusive focus on how well policies represent *the central tendency*, rather then the full distribution of preferences (c.f. Golder and Stramski, 2010). Studies of responsiveness evaluate the quality of representation based on the relationship between policy outcomes and the average preferences in states. Similarly, studies utilizing congruence base their normative claims on the match between policy outcomes and the *median* voter.[3] Essentially, both approaches treat states as "actors with distinct

[2]To anticipate our empirical application, consider the example of the minimum wage. The fact that a majority of Americans would prefer raising the federal minimum wage to $10.10 compared to the status quo of $7.25 does not imply that an average American would prefer the minimum wage to be set at $10.10.

[3]In the case where policies are binary, the majority opinion is simply the preference of the median voter.

preferences", and ignore the role of disagreement *within* states (though see Gerber and Lewis, 2004 and McCarty et al., 2015 for two exceptions).

This practice is problematic for two related reasons. First, to the extent that the variation of preferences *within* constituencies that are represented with a single policy is large, even if policies are close to what is preferred by the average citizen, they could be far from the preferences of most citizens, a case described by Figure 1c. For instance, if a state consists of localities that either prefer very low taxes or very high taxes, a moderate tax rate will represent preferences less well than in states where every locality prefers a moderate tax rate, even if the average tax preferences is the same in both states. Put more simply, constituencies with diverse preferences are harder to represent with a single policy.

Second, an exclusive focus on average preferences within states precludes researchers from exploring the ideal level of decentralization. In the case of issues for which preferences vary greatly within states, it is possible that policies should be delegated to lower levels of government (e.g. cities or school districts). Conversely, if preferences are very similar across states, the delegation of policy-making at the level of states may not produce efficiency gains, compared to a single national policy. Again, while the theoretical distinction between "representing the average voter" and "representing voters, on average" has been noted as early as in Achen (1978) this notion has not been appreciated in empirical research on the opinion-policy relationship.

A distance-based approach to measure of policy representation

Our approach to address the issues described above is based on the introduction of an alternative measure of representation that is based on the ideological distance between individual preferences and policies. Such measures have been introduced to the literature by Christopher Achen (1978) and have been successfully applied to the study of dyadic representation (e.g. Bafumi and Herron, 2010; Jessee, 2009). At the same time, this approach has not been utilized in the context of the opinion-policy relationship.

The key assumption behind distance-based measures of representation is that there exists an ideological space on which preferences and policies can be directly compared. To formalize this idea, denote policies enacted in a set of jurisdictions (i.e. states) by $W_s \in \mathbb{R}^+$ and denote the preference of an individual i residing in that jurisdiction by $\theta_{i,s}$. In this framework, Achen (1978) proposes two measures to evaluate representation. First, he defines *centrism* as the squared distance between average preference and policy ($C_s = (W_s - E[\theta_{i,s} \mid s])^2$). Second, he defines *proximity* as the average squared distance between individual preferences and policy ($P_s = E[(W_s - \theta_{i,s}) \mid s]^2$).[4]

While proximity and centrism represent a great conceptual advance in the study of representation, their applicability in our setting is limited by two issues. First, because the metric of both measures is the *squared* ideological distance between preferences and policies, they can only assess representa-

[4]Note, that in Achen (1978) the comparison is not between citizen preferences and policy but rather between the ideology of voters and their legislators; and that both centrism and proximity are defined only at the level of legislative districts. As it will be clear, this will be of great importance in our setting.

tion based on the *magnitude* but not the *direction* of the deviation between opinion and policy. Second, because Achen (1978) defines both proximity and centrism at a fixed the level of jurisdictions (his goal is to evaluate how well members of Congress represent their districts) these measures are not well suited to assess how the overall degree of representation would change due to changes in the level of jurisdictions at which policies are enacted in.

In order to overcome these issues, we introduce a new measure of representation, which we call *bias*.[5] We define bias as the signed deviation between an individual's preference and the policy enacted in the jurisdiction where she resides. Formally, $B_{i,s} = \theta_{i,s} - W_s$, so that positive values of bias imply policies that are more conservative than individual preferences. Aggregating bias across groups of individuals, we can estimate average bias within a a states or other jurisdictions, or even the the entire country (i.e. $B^{national} = E[\theta_{i,s} - W_s]$).

Given some estimates of policy bias, we can explore the quality of representation from three distinct angles. First, average bias, aggregated at any level of aggregation is informative of the direction and magnitude of the ideological slant of public policy, as faced by a certain group. Second, the variation of bias within a jurisdiction can be used to assess whether policy decisions are delegated to the right level: if the same policy is overly conservative compared to the preferences of some individuals in a jurisdiction but too liberal for others, than it is likely that the jurisdiction is just too heterogeneous to be served by a single policy. Third, exploring the

[5]We note here that Achen (1978) also uses the term "bias" to describe the intercept of the regression equation predicting legislator ideology with district average ideology.

correlates of bias (e.g. whether policy is closer to preferences in states with direct democratic institutions) can help us explain how the quality of representation could be improved.

Research design

The central goal of this paper is to introduce new approaches both in terms of data collection and statistical analysis in order to (1) obtain estimates of preferences and policy that are measured on the same scale, (2) explore the magnitude and variation of these preferences within and across states, and (3) use the resulting data to make evaluate about how well constituent preferences are represented. In this section, we first discuss minimum wage laws – the subject of our study, and then go on to describe our approach to estimating minimum wage preferences.

Minimum wage laws

Our empirical analysis compares citizen preferences about the minimum wage to corresponding state laws. This focus comes with three distinct advantages. First, this policy is extremely salient and it is "easy" enough to expect individuals to form meaningful opinions about it. The issue of raising the federal minimum wage was one of the cornerstone of the platforms of the Democratic presidential candidates in 2016, and local movements seeking to raise state minimum wages also got significant media coverage.

Second, minimum wage laws lend themselves well to our research design because we can map the universe of possible policies on to a meaningful scale (i.e. hourly wages in dollars). This property of minimum wage legislation

has been used in the past in to test theories of legislative behavior (Krehbiel and Rivers, 1988) as well as models of lawmaking (Clinton, 2012). Of course, minimum wage laws are complex: they can specify different rates for different groups (e.g. tipped workers) and can include provisions that require the indexation of the minimum wage to price indices. Throughout the paper, we maintain the assumption that both preferences and policies can be characterized by the highest minimum wage in a state, so that we can make meaningful comparisons between preferences and policies.[6]

A final advantage of our focus on minimum wage laws is that, because they are enacted both at the federal, state and local level, we can explore the role of policy delegation and sorting in the quality of representation. This issue has also been salient in the public discussions about minimum wage laws. As noted by Seltzer (1995), a key point in the debate over the Fair Labor Standards Act of 1938 was "whether to permit regional minimum wage differentials" (p. 1307). Senators from Southern states argued that "regional wage differentials were justified because of higher transportation costs and the lower cost of living in the South" (p. 1307) [7]. More recently, a number of states have passed so called "preemption" laws that bar cities and counties from setting minimum wage laws higher than that in their states. This multi-layered nature of minimum wage laws also implies that federal legislation could in principle dampen the responsiveness of state policies to mass preferences or that state laws could limit responsiveness at the local level.

[6] An additional complication is that variation in costs of living makes it difficult to compare policy bias across states. Thus, when interpreting such comparisons it should be noted that a dollar difference between preferred and actual minimum wages amounts to a larger bias in a cheaper locality.

[7] The minimum wage proposed by the FLSA was higher than the then prevailing wage in the South but below it in the North.

One possible objection for the use of minimum wage laws is that, given the complexity of the mechanisms through which they operate, individuals might simply be unable to form meaningful opinions about it or if they do, it is not possible to elicit such attitudes in with surveys. At the same time, even if citizens do not understand how minimum wages work, they can still evaluate how they impact the financial situation of their own and others (e.g. Fiorina, 1981) and thus form preferences over future policy changes based on their personal experience. Moreover, as demonstrated by recent research (Ansolabehere et al., 2014) survey respondents can handle questions involving numbers if they involve familiar quantities (such as gas prices) or if the question itself provides some anchor about the scale (e.g. the historical range of unemployment).

Measuring average preferences within states from survey data

We fielded a national survey to 3,500 respondents from YouGov's online panel in the winter of 2016.[8] The key innovation of the survey was that we measured preferences about the minimum wage using the following open-ended question: "The [respondent's state] minimum wage is $X an hour. How much (in dollars) do you think your state's minimum wage should be (0 meaning there should be none)?"[9] While we provide a detailed summary of this dataset in Online Appendix A, it is worth pointing out

[8]In order to approximate a representative sample of the adult population, YouGov employs matched sampling that involves taking a stratified random sample of the target population and then matching available internet respondents to the target sample. Such samples are shown to closely resemble the correlational structure of random samples using telephones and residential addresses (Ansolabehere and Schaffner, 2014).

[9]The survey was programmed so that each respondent was provided the actual minimum wage in his or her state. While we chose to provide "anchors" to survey takers with the current status quo in order to reduce measurement error in these responses (c.f. Ansolabehere et al., 2014), such anchors could potentially lead to the overestimation of responsiveness. To assess this possibility, we repeat our analysis of responsiveness using the responses to a similar question about the federal minimum wage (for which the anchor was identical across states), yielding very similar results.

that the range of survey responses (about 95% between 0 and $15.00) suggests that most survey takers understood the question and provided meaningful responses.[10]

The key challenge to this approach is to estimate state-specific average opinion from individual data, even though our survey is not representative at the level of states and contains only a handful of interviews in some of the smallest states. Following recent literature (e.g. Lax and Phillips, 2009; Ghitza and Gelman, 2013) we utilize multilevel regression and post-stratification (MrP), an approach that has been shown to produce more accurate estimates of state-level opinion than traditional techniques (Lax and Phillips, 2009, 2012). In particular, we model the preferred level of the state minimum wage expressed by each survey taker using demographic variables (age, gender, race, income, and education) as well as state-specific random intercepts. We then use Census data to post-stratify predicted mean preferences within each cell to obtain estimates of average state preferences. We provide additional details about this procedure in Online Appendix A.

Measuring the variation of preferences within states from referenda results

In order to generate localized estimates of minimum wage preferences, we use the results of six ballot initiatives that took place in Arizona, Colorado, Maine and Washington between 2013 and 2016, each seeking to raise the state or local minimum wages. (see Table 1 for a summary)[11] We generate

[10] We top-coded these responses at $25.00 because for a handful of responses (e.g. 1133 or 810) it seemed likely that they reflect typing mistakes rather than genuine preferences. This decision does not affect our results because it affects less than 0.5% of our data.

[11] This data is freely available online in spreadsheet forms at Arizona, Colorado, Maine and Washington. Because states vary in the level of aggregation at which they report election results (counties in Colorado

three sets of estimates from this data: first, relying on the results of the four statewide referenda we generate county level estimates of minimum wage preferences in each state. Second, relying on the three referenda taking place in Washington we generate local preference estimates in that state at the level of cities. Finally, exploiting the tight correlation between the resulting preference estimates and city-level estimates of ideology provided in Tausanovich and Warshaw (2014) we also extrapolate preferences in 20 cities with local minimum wage laws.

Table 1: Minimum wage ballot initiatives in 2016

Initiative	Jurisdiction	Type	Date	Status quo	Proposal
Proposition 206	Arizona	Statewide	2016	8.15	10
Amendment 70	Colorado	Statewide	2016	8.56	9.3
Question 4	Maine	Statewide	2016	7.5	9
Initiative 1433	Washington	Statewide	2016	9.55	11
Initiative No. 1	Tacoma	City	2015	9.55	10.35
Initiative No. 1B	Tacoma	City	2015	10.35	15
Proposition 1	Seatac	City	2013	9.32	15

Our empirical approach builds on the group-based IRT model discussed in Caughey and Warshaw (2015) with the modification that in this case the difficulty parameters are known. Following the literature on spatial models of voting (e.g. Clinton et al., 2004; Jessee, 2009) we assume that individual support for a ballot initiative is governed by the following random utility model:

$$P(v_{i,c} = 1) = P\left(\beta\left((\theta_{i,c} - w_c)^2 - (\theta_{i,c} - W_c)^2\right) > \epsilon_{i,c}\right), \tag{1}$$

and Arizona, and precincts in Maine and Washington) we present our results at the county level, the smallest units for which vote counts are available from each of the four state.

where $v_{i,c} = \{0, 1\}$ is an indicator variable that takes on 1 if an individual cast a supporting vote, θ is an individual's ideal minimum wage, w_c and W_c denote the status quo and the proposed minimum wage in county c, respectively. Moreover, β denotes a discrimination parameter, measuring the extent to which voting is influenced by the relative utility of policies vis-a-vis a stochastic shock and $\epsilon_{i,c}$ are IID stochastic disturbances distributed as standard normal. Combining this model with the assumption that individual preferences are distributed normally within each county, so that $\theta_{i,c} \sim N(\Theta_c, \sigma)$.[12] we can establish the relationship between the moments of the distribution of preferences in each county and the observed vote shares (denoted by $V_c = E(v_{i,c} \mid c = c)$) (Mislevy, 1983, p. 278; Caughey and Warshaw, 2015, p. 200).

$$E(V_c) = \Phi \left(\frac{\Theta_c - \frac{W_c + w_c}{2}}{\sqrt{\frac{1}{4\beta^2 (W_c - w_c)^2} + \sigma^2}} \right) \tag{2}$$

We estimate the parameters of this model (i.e. Θ_c, σ and β) using a fully Bayesian procedure in which we first specify prior distributions for each parameter of interest.[13]. Then, we use Monte Carlo simulation to obtain draws from the posterior distribution over the model's parameters by sampling from the conditional posterior distributions of each parameter given the current simulated values of all other parameters (Jessee, 2009). The resulting series of random draws from the posterior distribution over

[12] In this specification we implicitly assume that the variance of preferences is identical across counties. In Appendix B we show that our results remain virtually identical when we relax this assumption.

[13] Following Caughey and Warshaw (2015) we use half-Cauchy priors with a mean of 0 and a scale of 2.5 for the two variance parameters (i.e. $Var(\Theta_c)$ and $Var(\theta_i|c)$); a lognormal prior for the discrimination parameter β. Finally, based on the results of Study 1 we use a weakly informative prior for the grand-mean of average preferences (i.e. $\Theta = E(\Theta_c)$) as $\Theta \sim N(10, 4)$.

the model's parameters are then used to make inferences. We provide diagnostics of the estimation procedure in Online Appendix B.

The intuition behind this approach is simple. First, within each state, the estimation procedure exploits the variation of election results across counties to estimate the variation of mean preferences. Second, the standard functional assumptions made about individual utility functions as well as the distribution of ideal points allow the procedure to "bridge" across states based on the content of the ballot initiatives. For instance, the same election results in two counties in Maine and Washington is indicative of a higher minimum wage preference in the latter, given that the cut-point between the status-quo and and the proposal was higher in Washington.

Results

We present our empirical results in two sections, each corresponding to one of of our main arguments. First, relying on our survey data on individual preferences and state-level estimates we assess the magnitude of policy bias, explore whose preferences such bias policies reflect the most and finally discuss how direct democratic institutions can mitigate the observed bias.. Second, relying on our local preference estimates generated from referenda results, we set out to assess the variability of preferences within states, explore the determinants of the observed variation and finally describe how decentralization can mitigate the bias resulting from overly heterogeneous jurisdictions.

The bias of public policy

How biased is public policy? Figure 2 plots the estimated average preference about the minimum wage against the actual state minimum wages that were in place at the time of the survey. A striking feature of the relationship between preferences and policy across the American states is that in each of the 50 states and the District of Columbia, the actual minimum wage is less than what the average voter would have preferred. In the median state in terms of bias (Texas), the deviation between preferences and policy is $2.30. In 40 states, the bias is between 1 and 3 dollars. We report estimates for each state in Online Appendix A.

Figure 2: Responsiveness and bias in state policies

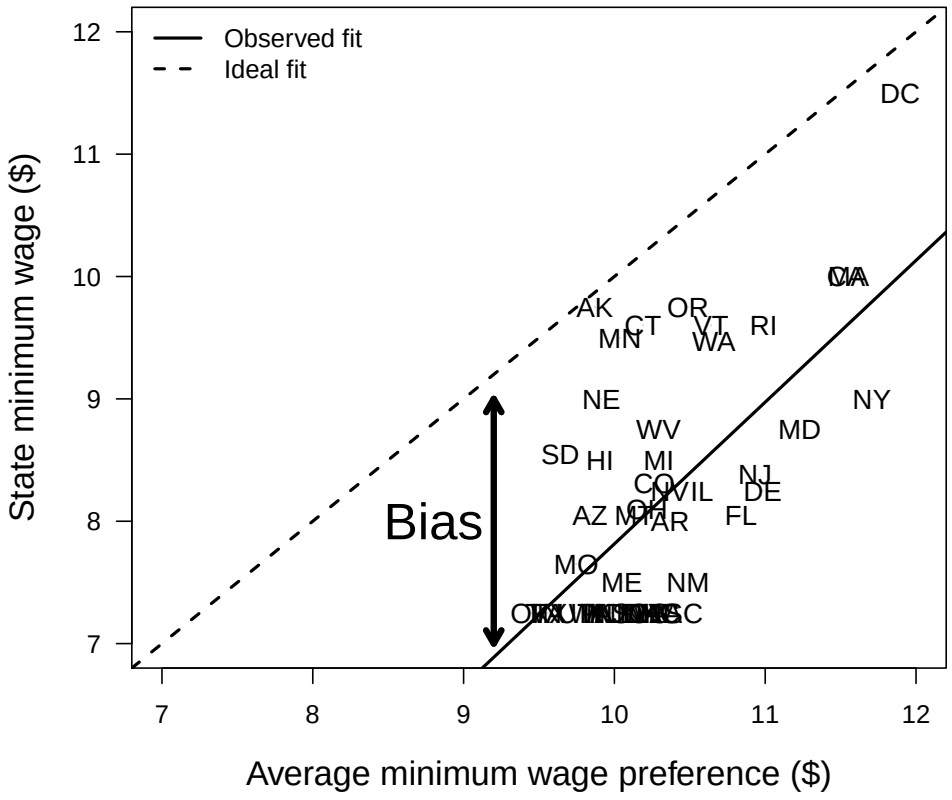

Note: The figure plots minimum wage laws in U.S. states (y-axis) against estimated average preferences (x-axis). The dashed line indicates the ideal relationship (i.e. policy corresponding to average preferences) so that observations below the line exhibit a conservative bias. The solid line depicts the estimated relationship between average opinion and policy using a bivariate regression of policy on MrP estimates of state opinion.

A good way to assess the magnitude of this bias is to compare it to the variation of average preferences across states. According to our estimates, the state with the most conservative minimum wage preferences was Oklahoma ($9.44) and the most liberal was New York ($11.70). The finding that policy bias within states was roughly the same as the full range of average preferences across states implies that even if more liberal states end up with more liberal policies is relatively inconsequential: liberal and conservative states alike have policies in place which are very conservative compared to average preferences in them.

Figure 2 also reveals how standard approaches to the opinion-policy relationship can reach misleading conclusions by restricting their attention to responsiveness, the association of policies and average preferences across states. In this instance, the slope of the linear relationship between average preferences and policy (the standard measure of responsiveness) is $\beta = 1.15$, with a 95% confidence interval of 0.74 to 1.55. This implies that a state where the average voter wants a $1 increase in the minimum wage (the difference between Texas, with an average preference of $9.55, and Illinois, at $10.53) is expected to have a policy that is more generous by about a dollar. Assuming that some states would enact minimum wages below $7.25 in the absence of federal law, the existence of the federal minimum wage leads to both *less* responsiveness and *less* bias.

Whom do biased policies favor? Given that one of the main justifications for minimum wages is to reduce poverty, a natural question to ask is whether the observed conservative bias reflects the lesser representation of the preferences of poor people (Bartels, 2008; Gilens, 2012). While the no-

tion that economic policy making in the U.S. disproportionally favors the financial interests of the affluent is widespread (e.g. Bonica et al., 2013), it has proved difficult to show that policy better represents the *preferences* of the rich. Rich and poor Americans often support the same policy alternative on many issues (Erikson, 2015) and it is not possible to recover relative bias from relative support for policies without measures based on the same scale.[14] As a result, there remains some disagreement on whether some of the conservative policy outcomes take place against the wishes of the poor (Gilens, 2012; Flavin, 2015) or instead, because poor people support them against their financial interest (e.g. Bartels, 2005; Kuziemko et al., 2015).

To explore the hypothesis that minimum wage laws in the states better reflect the preferences of the affluent than those of the poor, we simply compare the average deviation of state policy from preferences across income groups. Figure 3 provides striking evidence of the unequal representation of poor and affluent Americans: the policy bias faced by people earning less than $10,000 per year (3.08, 95% CI [$2.45, $3.72]) is much greater than the average deviation between preferences and policy. And Figure 3 clearly shows that policy bias decreases more or less monotonically by income, with the richest individuals getting policies that are on average 6 times closer to their preferences than do the poorest individuals.

[14]For instance, if we found that 40% of rich people and 60% of poor people would support raising the federal minimum wage from $7.25 to $12.00, we would not know the relative bias of the status quo from the average preference among the rich and the poor.

Figure 3: Policy bias by income

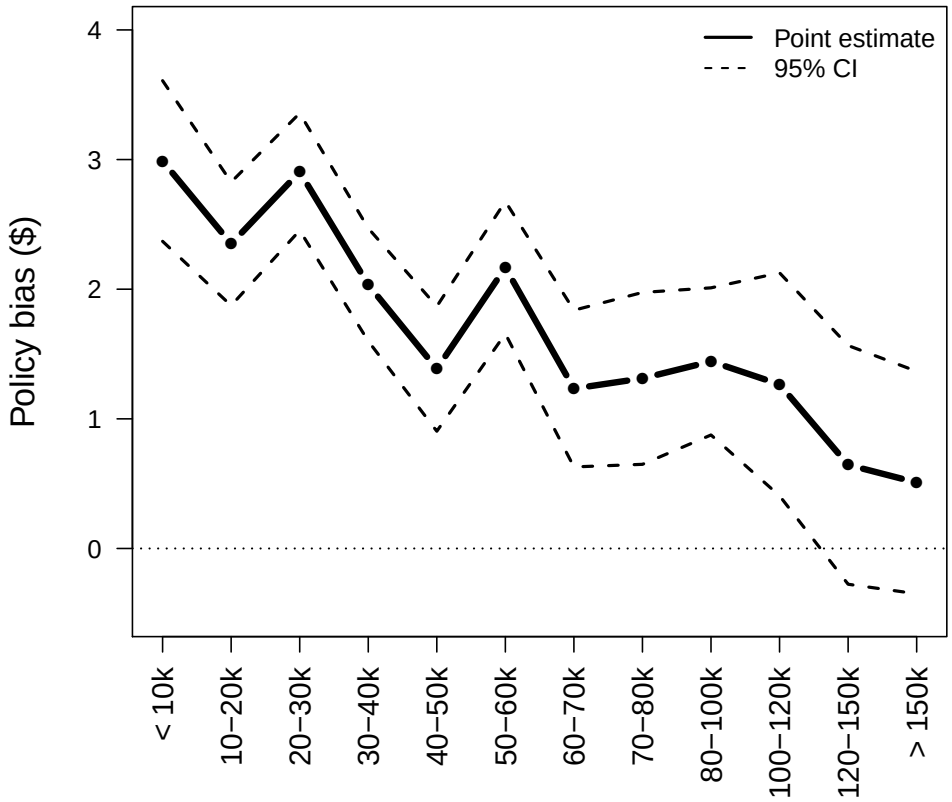

Note: The figure plots the average policy bias experienced by each income group ($E[\bar{w_{i,s}} - p_s|income]$). The dashed lines depict 95% confidence intervals. We weight the data using sampling weights.

How to reduce bias? Given that changes in state minimum wages have been brought about through popular initiatives in so many states (between 2006 and 2016, 16 referenda were held on minimum wage increases) a natural question to ask is whether access to such direct democratic institutions lowers policy bias. While the claim that access to direct democracy leads to policies that better reflect mass preferences has been made by many (e.g. Matsusaka, 2008; Leemann and Wasserfallen, 2016) studies that measure opinion and policies on different scales cannot directly compare policy bias across states.

In contrast to these earlier approaches however, our research design affords

a unique opportunity to quantify the impact of direct democracy on the magnitude of policy bias. First, based on our estimates of average preferences in each state we compared the deviation of preferences and policies across states with and without access to ballot initiatives. Policy bias was about \$2.59 in the median state without access to direct democracy in contrast to \$2.09 in the median state with direct democracy. Second, our county level estimates demonstrate that each of the four referenda that took place in 2016 reduced policy bias.

Taken together, these results point to both the promises and the limitations of direct democracy. First, it appears that successful ballot initiatives can indeed move policy closer to average preferences. Second, simply because minimum wage referenda lead to so large policy changes, it seems likely that the threat of popular initiative is insufficient to force legislatures to enact unbiased policies. Finally, the fact that even policies resulting from these successful referenda remained biased shows that elites can insulate policy outcomes from average preferences through the choice of policy proposals appearing on the ballot.

Delegation and the distribution of bias within states

The variation of policy preferences within states In contrast to existing approaches (though see Tausanovich and Warshaw, 2014), our research design also permits the description of the opinion-policy relationship below the level of states. As we have explained above, exploring the full distribution of preferences within states can qualify our conclusions about bias. To the extent that policy bias is homogeneous within states, shifting state policies is effective in improving representation. However, if

an average conservative bias within a state masks a diverse set of counties – some with more liberal and some with more conservative preferences than the prevailing state policy – then meaningfully enhancing representation is only possible by decentralizing policies further.

Figure 4 visualizes the variation of preferences within each of the four states. Similarly to Figure 1b, we also indicate the average preference within states (dashed lines) as well as the policies before (dotted lines) and after the referenda (straight lines). Two clear patterns emerge: first, status quo policies were more conservative than average preferences in each of the four states both before and after the ballot initiatives. While ballot initiatives reduced bias in each state substantially (by a minimum of 48% in Colorado to a maximum of 72% in Maine), policy still remained more conservative than the average preferences.

Second, the comparison of preferences and policy at such a low aggregation also reveals how within-state heterogeneity limits the representation of preferences. While policy bias at the aggregate was found to be conservative, there were counties in each of the four states where minimum wages were "too high" compared to average preferences even before they were increased. Similarly, even after the policy changes resulting from the referenda, there remained counties in each state where minimum wages still remained below what the average voter preferred.

Figure 4: The distribution of policy bias across counties

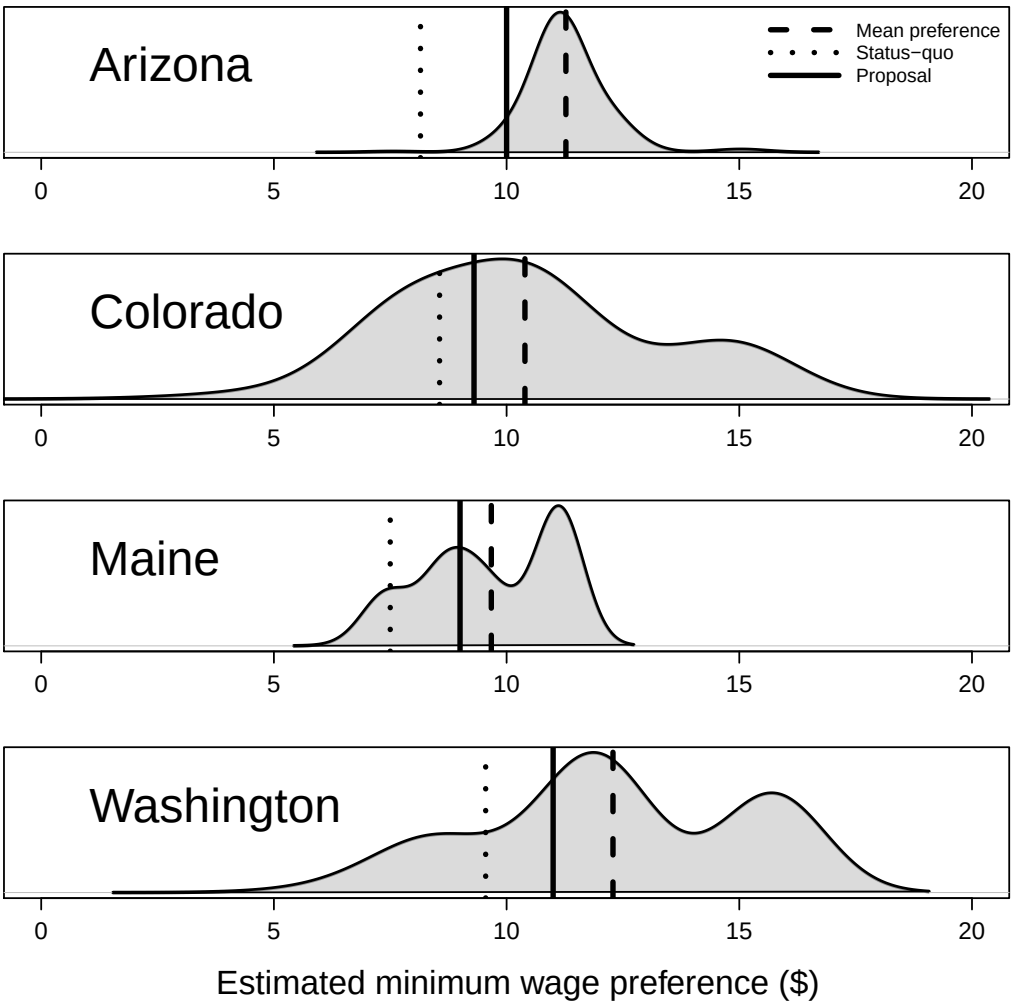

Note: Each plot visualizes the estimated distribution of minimum wage preferences within states, approximated by kernel density estimates, where each county is weighted according to the total number of votes cast in the 2016 minimum wage referenda.

The determinants of local policy bias Why are policy outcomes too liberal compared to preferences in some parts of a state but too conservative in others? Following recent research (Goldstein and You, 2017) we hypothesize that policy outcomes will be particularly biased in parts of a state which are different from the state as a whole in terms of policy preferences. In the context of minimum wage policies, we expect that counties that differ from their home state in terms of costs-of-living, or partisan preferences will be served less well by policy. In Figure 5 we thus com-

pare policy bias by relative costs of living (i.e. how expansive is a county compared to the state) as well as by relative partisanship (i.e. democratic vote share in the county compared to state results).

Figure 5: The distribution of policy bias across counties

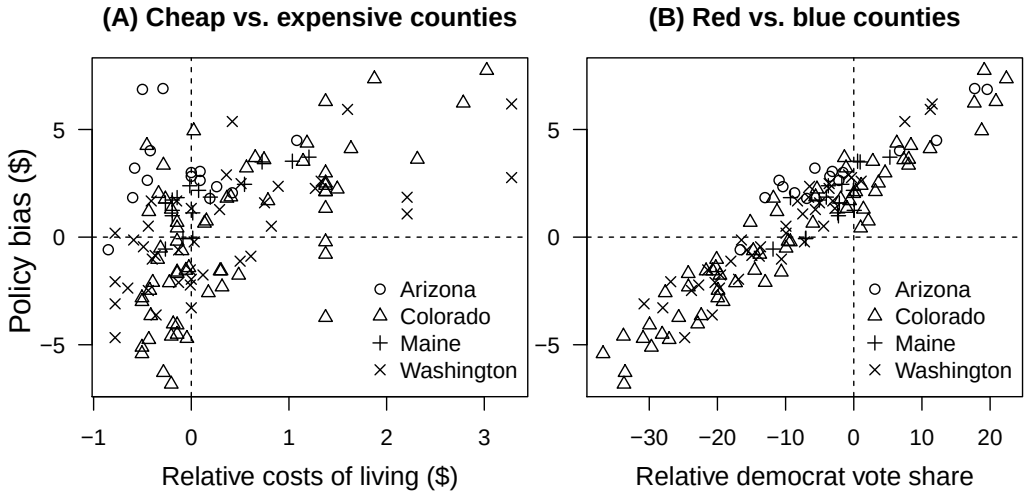

Note: The figures plot estimated policy-bias against relative costs of living (left panel) and relative democratic vote-share (right panel) in each county in the four states that held minimum wage referenda in 2016. Policy bias is defined as the difference between county level preference and state policy. Relative cost of living is defined as the difference between the estimated living wage in a county and its state. Relative democratic vote share is the difference between the share of two-party Obama votes in 2012 in the county and the state.

Our results show a tight correlation between policy bias and both predictor variables. On one hand, the left panel of the figure shows that counties that are more expensive than the state average experience greater conservative policy bias, while many counties that are cheaper than their state face too high minimum wages compared to citizen preferences. On the other hand, we find a similar – and even more pronounced– pattern in the case of partisan composition. Counties that voted Democrat in greater proportions in the 2012 Presidential election compared to their state's average face much greater conservative policy bias compared to relatively more conservative counties .

We note two additional patterns emerging from this exercise. First, even though there is a roughly linear relationship between relative costs of living and relative partisanship on one hand and policy bias on the other, some states which are cheaper or less Democratic leaning than others still experience policies that are too conservative. The reason for this is simply that state policies on average are more conservative than preferences. Second, and relatedly, Figure 5 again demonstrates the enormous heterogeneity within states: even a relatively expensive and liberal state such as Washington includes counties that are cheap and conservative, which as a consequence experience liberal policy bias.

Where should policies be delegated? The debate on whether state-house democracy leads to policy outcomes that represent mass preferences well is tightly connected to the debate on whether policies should be delegated to states in the first place (e.g. Lax and Phillips, 2009, p.382). To the extent that states differ from each other in terms of preferences, decentralization clearly enhances representation. However, as we have shown in the context of minimum wage laws, state policies deviate substantially from average preferences, and because state themselves are heterogeneous, most jurisdictions are not represented well by policy.

Two strategies have been proposed to mitigate this bias, pushing federalism in diametrically opposing directions. On one hand, proponents of raising the federal minimum wage would change policy outcomes in the states with the most conservative policies (i.e. the states that have no minimum wages) through centralization. On the other hand, local movements have successfully brought about minimum wage increases in cities

(such as San Francisco, CA; Flagstaff, AZ or Seattle, WA), essentially further decentralizing policy.

In order to assess the extent to which such changes in the level of policy making can enhance representation, we implemented the following procedure. First, exploiting three local and a statewide referenda in Washington and capitalizing on the availability of election results at the level of precincts we estimated city level minimum wage preferences in the 37 largest cities in Washington. Next, we merged the resulting data with city level ideology estimates produced by Tausanovich and Warhaw (2014) and generated predictions of minimum wage preferences in each of the 1500 cities in the US with a population over 25,000. Finally, we merged the resulting dataset with a comprehensive list of local minimum wage laws to estimate the impact of these local laws on policy bias.

The left panel of Figure 6 visualizes the distribution of policy bias in the 21 cities in our sample which had higher local minimum wages that their state and compare it with a counter-factual distribution of bias in the absence of these local laws. The dark-grey area depicts the actual distribution of bias (weighted by city-population) while the light-grey area shows the counter-factual distribution of bias in the case where state laws over-ride city minimum wages. The figure provides clear evidence that these local laws are efficient tools to reduce bias. The average distance between preferences and policies in that sample is \$1.73, while it would be about \$3.00 in the absence of city minimum wage laws, showing an effect size comparable to the estimated effect of direct democracy.

Figure 6: Policy bias with and without city minimum wages

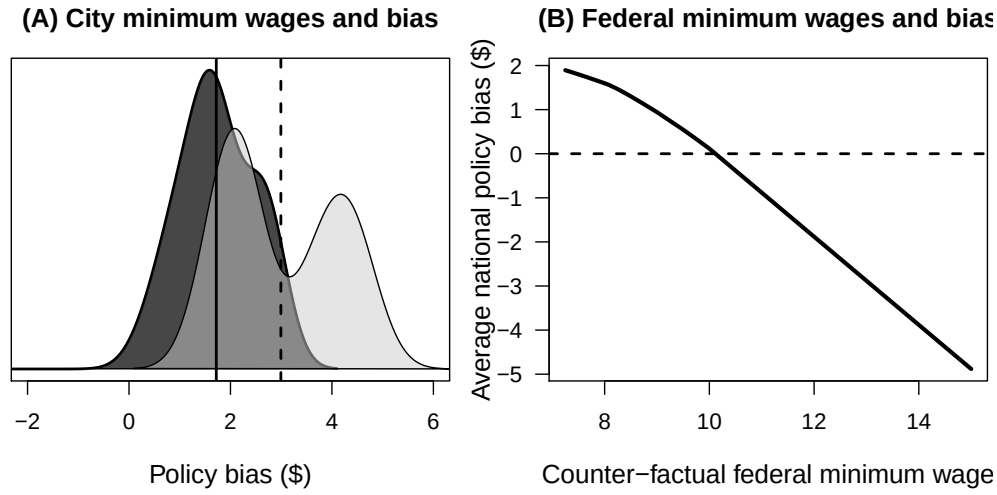

Panel A: The plot visualizes the distribution of bias across 21 cities with their own minimum wage laws (dark-grey area) and the hypothetical distribution of bias in the absence of local minimum wage laws (light area). Bias is the signed difference between preferences and policy. The distributions are approximated using kernel density estimates, where each city is weighted according to its population (as provided by Tausanovich and Warshaw, 2014). Panel B: The plot visualizes the estimated mean bias (averaged across respondents in our national survey data) for hypothetical values of the federal minimum wage.

We contrast these results with another set of counter-factual calculations in the right panel of Figure 6. In particular, we use our survey data to assess how a range of counter-factual federal minimum wage increases would change bias at the national level. The figure shows that moderate raises in the federal minimum wage that would increase the floor for state policies to the level of average preferences in the most conservative states would clearly enhance the degree of representation. For instance, according to our national survey data, increasing the federal minimum wage to $9.44, which is the average preference in Oklahoma, would reduce the average deviation of policy from opinion to below one dollar, amounting to a greater than 50% reduction in bias.

At the same time, the same results also make it clear that more aggressive

changes to the federal minimum wage (such as the proposed increase to $15.00) would essentially flip the sign of the average policy bias without reducing its magnitude. It is worth noting the stark contrast between this finding and possible conclusions drawn from using congruence. While recent surveys have repeatedly shown that a majority of Americans would support increasing the federal minimum wage to $15 (e.g. Pew, 2016), based on our results it seems that the support for such an extreme policy change is driven by an extremely conservative status quo rather than extremely liberal preferences.

Conclusion

What do our findings tell us about the quality of democratic government in the states? Based on our empirical results from this particular policy domain, two patterns stand out. First, even though statehouse democracy "works" in the sense that policy outcomes are tightly related to mass opinion across states, the same policies are biased in the sense that they are far removed from citizen preferences. Second, because the variation of preferences within states appears larger than across them we find that even if policies did reflect mass opinion on average, they would fall far from the preferences of most localities within a state.

Our analyses exploring the causes of policy bias highlight the role of direct democratic institutions. First, comparing policy bias across states we found that policy is substantially closer to average preferences in states with direct democratic institutions. Second, our analysis of the four ballot initiatives taking place in 2016 shows that the policy change resulting from these referenda greatly reduced bias. Third, direct democracy can

also ameliorate the democratic deficit resulting from within-state heterogeneity by making it possible for jurisdictions within states to enact their own policy.

Finally, our results also speak to the emerging literature on the unequal influence of the affluent in shaping policy. Given the continuing rise of economic inequality in the U.S., many observers claim that American public policy increasingly favors the interests of the affluent (e.g. Bonica et al., 2013). At the same time, because of the inferential challenges resulting from the incomparable scales on which preferences and policies are measured, the literature has struggled the show that policies better reflect the preferences of the rich than the poor. We believe that we are the first to show direct evidence showing that the anti-poor bias of a particular public policy, which takes place against the preferences of the least well-off rather than driven by some irrational demand from them for policies harming their interests.

Similar to many studies in the existing literature (e.g. Gerber, 1996; Lax and Philips, 2009; Lupia et al., 2010), our empirical analysis focuses on a single policy. As a consequence, our results do not necessary generalize to the opinion-policy relationship in the context of other issues. First, it is possible that in the case of other issues, policies are closer to the preferences of average citizens. This is likely to be the case for issues on which mass opinion is less divided simply because for such policies the incentives faced by pivotal legislators to please their constituents would be more aligned with majority support.

Second, it is possible that in the case of other policies the average bias

of policy points to the liberal rather than the conservative direction For instance, one property of minimum wage laws is that without indexation, policy drifts towards a conservative direction when no legislative action is taken. While this property describes many economic policy issues (e.g. thresholds for entitlements, regulations etc.) in certain issues areas where the status quo is closer to the liberal preferences, inertia may lead to liberal policy bias. For instance, since Roe v Wade, most legislation on abortion seeks to introduce regulations that would make abortion more difficult. To the extent that some of these initiatives are not introduced even though they enjoy majority support, policy would exhibit a liberal bias.

Third, our conclusions about the nature of geographic variation in preferences is likely to stand in the light of a large and growing body of empirical research on geographical sorting (e.g. Ansolabehere et al., 2006) finding scant evidence of spatial polarization in the U.S., at least not at the level of states. The notion that states are not too different from one another, but quite heterogeneous in terms of policy preferences can lead to two different conclusions. On one hand, for policies like the minimum wage it may seem reasonable to delegate policies to lower levels of government or craft legislation that explicitly differentiates laws at sub-state levels. On the other hand, the relative homogeneity of preferences across states might make the centralization of policy at the federal level more desirable as in the case of gay marriage.

At the same time, our results call into question the rather optimistic conclusions drawn by much of the existing literature. On one hand, we provide an example that very clearly shows that responsiveness does not imply

policies that are close to average preferences. While this point have been made by many (Achen, 1978; Lax and Phillips, 2012) researchers still use responsiveness to evaluate representation (e.g. Caughey and Warshaw, 2017). On the other hand, we also show that even policy change that is congruent with mass preferences does not necessarily result in policies that are close to average preferences.

Taken together it seems that for a richer understanding of statehouse democracy it is simply inevitable to develop research designs that permit the measurement of preferences and policy outcomes on the same scale. While this may prove more difficult in the case of policies which lack an intuitive natural metric to them (such as gun control or abortion) recent developments in the study of dyadic representation (Jessee, 2009; Bafumi and Herron, 2010) show great promise. New data and empirical approaches that allow the comparison of policy outcomes with preferences across a range of issues would open up new avenues for future research that could provide a more complete picture of representation.

References

Achen, C. H. 1978. Measuring representation. American Journal of Political Science, 475-510.

Ansolabehere, S, M. Meredith, and E. Snowberg. 2013. Asking about numbers: Why and how. Political Analysis 21(1) 48-69.

Ansolabehere, S. and Schaffner, B.F., 2014. Does survey mode still matter? Findings from a 2010 multi-mode comparison. Political Analysis, 22(3) 285-303.

Bafumi, J., and M. Herron. 2010.Leapfrog Representation and Extremism: A Study of American Voters and Their Members in Congress. American Political Science Review 104:519-42.

Bartels, Larry M. "Homer gets a tax cut: Inequality and public policy in the American mind." Perspectives on Politics (2005): 15-31.

Bartels L. 2008. Unequal Democracy: The Political Economy of the Gilded Age. Princeton, NJ: Princeton University Press

Bonica, Adam, Nolan McCarty, Keith T. Poole, and Howard Rosenthal. "Why hasn't democracy slowed rising inequality?." The Journal of Economic Perspectives 27, no. 3 (2013): 103-123.

Broockman, David E. 2016. Approaches to studying policy representation. Legislative Studies Quarterly 41(1): 181-215.

Card, David, and Alan B. Krueger. "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania." American Economic Review 84 (1994): 772-793.

Caughey, D. and Warshaw, C., 2014. Dynamic Representation in the American States, 1960-2012. *working paper*

Clinton, J. D. 2012. Congress, Lawmaking, and the Fair Labor Standards Act, 1971-2000. American Journal of Political Science 56(2): 355-372.

Downs, A., 1957. An economic theory of political action in a democracy. Journal of Political Economy, 65(2) 135-150.

Dube, Arindrajit, T. William Lester, and Michael Reich, Minimum Wage

Effects across State Borders: Estimates Using Contiguous Counties Review of Economics and Statistics, 92 (4), (2010), 945 964.

Erikson, R. S., G. C. Wright, and J. P. McIver. 1993 Statehouse democracy: Public opinion and policy in the American states Cambridge University Press

Erikson, R. S. 2015 Income Inequality and Policy Responsiveness Annual Review of Political Science 18:11-29

Fiorina, Morris P. "Retrospective voting in American national elections." (1981).

Flavin, Patrick. "Campaign finance laws, policy outcomes, and political equality in the American States." Political Research Quarterly 68, no. 1 (2015): 77-88.

Gerber, Elisabeth R. "Legislative response to the threat of popular initiatives." American Journal of Political Science (1996): 99-128.

Ghitza, Y. and Gelman, A., 2013. Deep interactions with MRP: Election turnout and voting patterns among small electoral subgroups. American Journal of Political Science, 57(3), pp.762-776.

Gilens, M. 2012. Affluence and Influence: Economic Inequality and Political Power in America. Princeton: Princeton University Press.

Gilens, M., and Page, B. I. 2014. Testing theories of American politics: Elites, interest groups, and average citizens. Perspectives on politics, 12(3), 564-581.

Glasmeier, A.K. 2017. Living wage calculator. Cambridge: Massachusetts Institute of Technology. http://livingwage.mit.edu.

Jessee, S., 2016. (How) Can We Estimate the Ideology of Citizens and Political Elites on the Same Scale?. American Journal of Political Science, 60(4), pp.1108-1124.

Krehbiel, K. and Rivers, D., 1988. The analysis of committee power: An application to Senate voting on the minimum wage. American Journal of Political Science, 32(4) 1151-1174.

Kuziemko, Ilyana, Michael I. Norton, Emmanuel Saez, and Stefanie Stantcheva. "How elastic are preferences for redistribution? Evidence from randomized survey experiments." The American Economic Review 105, no. 4 (2015): 1478-1508.

Lax, J. R., and Phillips, J. H. 2009. How should we estimate public opinion in the states? American Journal of Political Science 53(1): 107–121.

Lax, J. R., and Phillips, J. H. 2012. The democratic deficit in the states. American Journal of Political Science, 56(1), 148-166.

Leemann, Lucas and Wasserfallen, Fabio. 2016. The Democratic Effect of Direct Democracy. American Political Science Review 110(4): 1-13.

Matsusaka, John G. "Problems with a methodology used to evaluate the voter initiative." The Journal of Politics 63, no. 4 (2001): 1250-1256.

Matsusaka, John G. 2010. Popular control of public policy: A quantitative approach. Quarterly Journal of Political Science 5(2) 133-167.

Neumark, David, and William L. Wascher. Minimum wages. MIT Press, 2008.

Sabl, A. 2015. The Two Cultures of Democratic Theory: Responsiveness, democratic quality, and the empirical-normative divide. Perspectives on Politics 13(2): 345-365.

Seltzer, Andrew J. 1995. The Political Economy of the Fair Labor Standards Act of 1938, Journal of Political Economy 103 (6): 1302-42.

Tausanovitch, C. and Warshaw, C., 2014. Representation in Municipal Government. American Political Science Review, 108(03) 605-641.

Oates, William .E. (1999). An essay on fiscal federalism. Journal of Economic Literature, 37(3), pp.1120-1149.

Tiebout, Charles M.Measuring state tax progressivity (1956). A pure theory of local expenditures. The Journal of Political Economy, 416-424.