Stability and change in the opinion-policy relationship : Evidence from state minimum wage laws*

Gabor Simonovits[†]
Central European University

Alexander Bor[‡]
Aarhus University

December 20, 2021

Abstract

Recent studies have documented large discrepancies between mass preferences and policies in U.S. states consistent with theories that highlight the oversized influence of affluent Americans on policymaking. In this note, we replicate and extend a recent such study (Simonovits, Guess, and Nagler, 2019) to assess how policy bias evolves in time. Specifically, relying on novel data and methods, we construct measures of minimum wage preferences and compare them to observed policies in each state for the years of 2014, 2016, 2019, and 2021. We demonstrate that, averaged across states, policy change closely tracked a pronounced increase in preferences for higher minimum wages, but the size of policy bias remained relatively stable. However, this national pattern hides an increasingly polarized policy landscape: in many states, insufficient responsiveness led to an increasing deviation between preferences and policies, while in other states policy changes – larger than preference changes – closed initial policy bias.

Word count: 3,407

This working paper has not been peer-reviewed yet.

*We thank Andy Guess for his work on initial phases of his project. We are also grateful for comments from Neal Beck, Martin Vinæs Larsen, Michael Dorsch and seminar participants at Aarhus University and Oxford University.

[†]Assistant Professor, Department of Political Science, Central European University. Email: simonovits@ceu.edu

[‡]Postdoctoral Researcher, Department of Political Science, Aarhus University

Stability and change in the opinion-policy relationship: Evidence from minimum wage laws

To what extent do public policies reflect mass preferences? Some recent studies have documented large deviations between public opinion and public policies across a range of issue domains (Lax and Phillips, 2012) and others pointed to patterns consistent with the outsized influence of the wealthy and the power of interest groups (Gilens and Page, 2014; Simonovits, Guess, and Nagler, 2019). However to our knowledge all studies that collected the data necessary to compare policies and opinions on the same scale – a feature necessary to quantify policy bias – have used cross-sectional data. Thus, an important question left open in this literature is whether such “gaps of representation” are eventually closed as a result of legislators responding to mass opinion.

Providing evidence on the dynamic relationship between average preferences and policies is challenging for two separate reasons. First, in order to compare preferences and policies on the same scale one needs to collect data on the two which are measured on the same scale. Second, the units of analysis (i.e. states) need to be observed repeatedly so that dynamics can be studied. To our knowledge no such studies exist. Instead, studies have explored the related question of whether changes in preferences are *correlated* with corresponding policy-change (Caughey and Warshaw, 2018). However, policy could move in the same direction with mass opinion *and* stay far removed from it so long as initial divergence is large and/or the degree of dynamic responsiveness is insufficient. Moreover, a focus on responsiveness focuses on the *average* relationship between preferences and policies, and as such using it as an evaluative standard is likely to mask important heterogeneities across states.

In this paper we deploy a research design and marshal data that allows us to precisely describe the dynamic relationship between mass opinion and policy - at least in one specific issue domain. We replicate and expand a recent study by Simonovits, Guess, and Nagler (2019) to explore if its conclusion – a pronounced and uniform conservative bias – has stood. Our analysis relies on seven national surveys fielded between 2014 and 2021 that measure preferences about the minimum wage in U.S. dollars. Using an augmented Bayesian dynamic multilevel model with post-stratification (MRP), we obtain estimates of average preferences for each U.S. state in the years 2014, 2016, 2019 and 2021 and match the resulting data with minimum wage laws in the same years. As

we measure preferences and and policies on the same scale we can assess the dynamics of representation in each state simply by tracking policy bias (the deviation between preferences and policies through time).

We report two key findings. First, on average the representation of mass preferences in minimum wages has stayed quite stable across the 7 years we study. Averaged across states, the difference between average preferences and policies hovered between 2 and 3 dollars. Such stability is quite remarkable given the large number of legislative changes¹ as well as the pronounced upward trend in mass preferences, visible in most states. Second, considering variation across states, we find a great degree of heterogeneity despite the short-span of our study: in many states, legislative inaction together with increasing support for higher minimum wages has lead to significantly *increasing* gaps between preferences and policies. In contrast, other states have seen overly responsive policy-making that led to a significant improvement of representation, at least on the short run. Taken together, our findings point to the importance of using measures of representation that can uncover important differences across states.

Taken together these findings showcase the importance of replicating descriptive findings in political economy. In the absence of such replications, there is no way for us to know if classic results of representation still characterize well the empirical relationship between public opinion and public policy. Insofar as we care not only about qualitative (e.g. preferences and policy outcomes are correlated) but also quantitative patterns (e.g. minimum wages are about \$ 2 lower than preferences) this issue is all the more acute. Our approach reveals important changes in the opinion policy relationship despite studying it within a relatively short time-period suggesting that similar replications could overturn conventional wisdom across a wider range of contexts.

Research design

Our empirical analysis explores the relationship between average preferences about the minimum wage and corresponding state laws between 2014 and 2021. Our focus on a single policy issue – rather than a composite measure of policies – spanning a relatively short time period come with the cost of limited generalizability; how much and how fast policies respond to changes in public opinion might be different in other periods or for other policies.

¹Ballot initiatives alone led to 12 minimum wage increases in the studied period.

At the same time, our design also comes with distinct advantages. First, the minimum wage is a highly salient issue, which – we believe – may be prone to short-term changes due to changes in consumer prices. For example, if a citizen notices that rents went up substantially in their hometown, they could plausibly update their belief whether it’s possible to make ends meet on the current minimum wage. Second, minimum wage laws are changed relatively frequently and easily. Case in point, during the period we study 13 statewide ballot initiatives have been voted on. In short, minimum wages constitute a most likely case for policies that closely track changes in public opinion. Finally, quantifying representation on this issue is easy as changes in policies and preferences can be measured on a meaningful and comparable cardinal scale.

Data

We rely on seven national surveys where respondents were asked to express their preferred levels of the minimum wage in dollars. We report detailed information about the datasets – including question wording, summary statistics and minor coding decisions – in Section A in the Online Appendix (OA). In the “supersurvey” we created from these data, we have 21,951 observations from 15,635 unique respondents. We use the following demographic covariates in our analyses: age (4 categories), sex (2 categories), education (5 categories), race/ethnicity (2 categories), state of residence (51 categories including Washington DC). Thereby, our models identify $4 \times 2 \times 5 \times 2 \times 51 = 4080$ types of US citizens. Our models also include year of survey (4 categories), and an interaction between state and year (204 categories). As a state-year level smoother, we rely on Presidential election results using simple linear intrapolation to estimate the state-level two-party vote share of the Democratic party in the November prior to data collection. For example, while for the 2021 data, we rely on the 2020 November vote, for 2014 we estimate 2013 vote as the weighted average of 2010 and 2014 elections (with weights of 0.25 and 0.75, respectively).

Note that the seven surveys diverge in terms of size, mode and sampling. In particular, five of the seven surveys rely on online opt-in panels. However, because we rely on multilevel modeling and post-stratification, our state-level estimates of public opinion will approximate true state-level preferences as long as inclusion to these samples is ignorable to our dependent variable – i.e. minimum wage preferences – conditional on the demographic variables we include in our models.

For the second stage of our estimation, we rely on US Census microdata from the Amer-

ican Community Surveys (ACS) (Ruggles et al., 2020). We use 5-year samples for each year prior to our data collection when it is available (thus 2013, 2015, and 2018). For the 2021 data, we rely on the most recent 5-year ACS available, i.e. 2019 data. We aggregate these microdata to compute the frequency of each category of citizens.

Estimating state-level preferences

To obtain state-specific preferences for each of the four years we study, we follow recent literature (Ghitza and Gelman, 2013; Lax and Phillips, 2012) and utilize multilevel regression and poststratification (MrP). Our model is non-standard in that we predict opinion on a cardinal scale, and we are among the first few in the literature to explore *changes* in public opinion (Gelman et al., 2016). Using cardinal minimum wage preferences poses an important challenge as the distribution of preferences is bi-modal. The majority of preferences is centered around \$11, but a minority of respondents would abolish state minimum wage policies and therefore indicate a preference of \$0. Simple linear models are unable to account for this distribution underestimating the frequency of \$0 preferences and provide overly confident estimates. Instead, we build our model in two steps. First, a multilevel logistic regression predicts whether participants want a state minimum wage or not (i.e. prefer \$0 or not). Second, a multilevel linear regression predicts the level of minimum wage preference among respondents who want to have (a non-zero) state minimum wage ($N = 19,860$). We build both models in a step-wise procedure, relying on leave-on-out (LOO) cross-validation to ensure that additional parameters improve out-of-sample predictions, and thus avoid over-fitting. Figure OA1 displays the coefficient estimates from these two models. Section C.3 in the OA demonstrates that combining two regression models not only improves model fit, but also has important substantive implications as a naive linear regression model underestimates preferences and thus policy bias, and also the uncertainty in our predictions.

Both logistic and linear models account for the temporal dynamics of public opinion by including year and state-year variables. Our approach is superior to running separate models for each year because it allows partial pooling across time. This step is also crucial for estimating state-level minimum wage preferences from surveys with only federal minimum wage preferences. Moreover, Democratic vote-share also varies across both states and years, and thus we can use it to pool information about time trends.

Finally, our models include a random intercept representing responses to questions on

the *state* versus the *federal* minimum wage. In the two surveys where we measured both preferences (for the years 2014 and 2016) we use both responses of each respondent and the random intercept estimates the average difference between state and federal minimum wage preferences. We then use this estimate to extrapolate state-level preferences in the 2014 surveys. We are confident that this extrapolation yields valid estimates for two reasons. First, respondents overwhelmingly indicated similar preferences for federal and state-level minimum wages, evident in the fact that the coefficient estimate for the item was 0 in both models. Second, we explored if the gap in federal and state-level preferences is different across states, but found that accounting for this potential heterogeneity does not improve our model fit.

We build our multilevel models in RStanArm, an R library implementing Stan software’s full Bayesian statistical inference with MCMC sampling. Bayesian analysis has two main advantages: first, weakly informative priors help even very complex multilevel models (like ours) to converge. Second, the simulations result in posterior draws for each of our parameters and thus we can also obtain posterior distributions for minimum wage preference for all possible combinations of our independent variables (demographic groups, states and years). In essence, our analysis yields two posterior distributions for all 4,080 types of US citizens in each of our 4 years: the share of respondents within the category whose preference is \$0 and the average preference for those whose preference is non-0. We multiply these two posterior distributions to arrive at a combined minimum wage preference estimate. Finally, these estimates are weighted (post-stratified) to the US Census to account for the unique demographic composition of each state in each year and are aggregated to state-years to obtain our MrP estimates of average preferences in each state and year pair. We describe our model specification as well as our priors in Section B.1 in the OA. In Section C.2 of the OA, we present evidence of the validity of our state-level estimates. Both across and within states, we observe a strong correlation between our estimates and rental prices.

Our goal is to explore the evolution and dynamic relationship between public opinion about the minimum wage and corresponding policies in each state. To do so, we use the repeated measures obtained from the MrP procedure and merge the panel dataset of average preferences with corresponding data on the minimum wage laws that were in place in each U.S. state in the years 2014, 2016, 2019, and 2021. Based on these data we

can compute *policy bias* – the deviation between average preferences and policies in each state (Simonovits, Guess, and Nagler, 2019).

Results

We first consider the relationship between preferences and policies in each year included in our study. The four panels in Figure 1 plot state policies in 2014, 2016, 2019 and 2021 against average preferences in the corresponding states and years. Three clear patterns stand out. First, in each year, we observe a substantive conservative policy bias; almost every point in our figure is below the diagonal. Second, in each of the four years, policies were strongly related to variation in preferences; states with preferences for higher minimum wages had on average more generous policies.² Finally, the relationship between states and preferences appears to be strengthening in two ways: the minimum wage difference associated to a \$1 difference in preferences is increasing, from 74 cents in 2014 to \$2.26 in 2021. Besides, the explained variation in policies by preferences is also steadily increasing from an R-squared of 0.48 in 2014 to 0.61 in 2021.

²See Table OA3 for yearly estimates of static responsiveness. Note that we propagate uncertainty in preferences using the method of composition Caughey and Warshaw (2015)

Figure 1: Static policy responsiveness by year

Note: Each state is denoted with a point. Black lines represent the diagonals. Blue lines are linear trends.

While these patterns suggest that *on average* state-house democracy has improved in this time period, we continue our analyses by disaggregating temporal variation in responsiveness state-by-state. Figure 2 plots our estimates of average preferences (x-axis) against the policies for each US state. Vertical bars around each point denote 95% highest density continuous intervals. More biased policies (those further from the diagonal denoting congruence between preferences and policies) are denoted with darker shades of gray.

Figure 2: The trajectory of U.S. minimum wage preferences and policies between 2014 and 2021

A number of notable lessons emerge. First, preferences are monotonically increasing in almost all states, on average by \$2.3. We observe the largest increase in Massachusetts (\$3.2) and the smallest in Oklahoma (\$1.5). Second, the median minimum wage across states also steadily increased (medians: 2013–\$7.25; 2016–\$8.05; 2019–\$8.5; 2021–\$9.25) on average by a similar amount than preferences (mean change = \$1.8), yet this average change masks a much larger heterogeneity across states: in twenty states – notably in Southern US – policy does not change at all, either for the lack of a state minimum wages (e.g. Alabama, Louisiana) or for a reluctance of state legislatures to change them (e.g. Texas, Utah). In other states (e.g. in Virginia, Illinois), policy is changed only once in the 7 years of our study period, thereby creating an ebb and flow in policy bias. Finally, a third group of states has updated policies regularly, yielding a relatively stable

policy bias (e.g. Oregon, Connecticut). In short, Figure 2 warns us that it is not easy to summarise the diversity of preference-policy trajectories with just a single estimate of policy responsiveness.

Figure 3 summarizes these patterns by exploring temporal variation in policy bias. It shows that the median bias hovered between \$2 and \$3 throughout our study period: it grew from \$2.2 in 2014 to \$3.1 in 2016, then decreased to \$2.5 in 2019, and grew again to \$2.8 in 2021. Meanwhile, policy bias in minimum wage laws also reflects an increasing polarization in the US. The range of biases grew steadily from \$2.7 in 2013 to \$3.3 in 2016, \$4.9 in 2019 and \$6.1 in 2021. No state had a liberal bias in 2014, and the smallest conservative bias was in Oregon at \$1. By 2021, ten states had less conservative bias than that, and three states (California, Washington, and DC) had a liberal policy bias.

Figure 3: The change in policy bias by year. While the mean bias remains similar, there are increasing differences across states

Note: Each state is denoted with thin line. The thick line is the national average policy bias weighted by state population size. The error bars denote 67%, 89% and 97% credible intervals, respectively.

Discussion

Our study offers new insights on the dynamic relationship between policies and public opinion in the United States and contributes to the rapidly evolving methodological arsenal for studying it. Going beyond existing research (Caughey and Warshaw, 2018), we capitalize our unique data and design to quantify the extent to which public policies

track changes in preferences. Moreover, our ability to measure preferences and policies on the same continuous scale allows us to investigate how states differ in translating changes in preferences in corresponding policies.

Based on our analysis of minimum wage laws, it appears that statehouse democracy “works” on average in the sense that changes in policies are related to fluctuations of mass opinion within states. At the same time, we find that even though on average, policy changes followed public opinion, such responsiveness was insufficient to close a large gap between the minimum wages citizens want and what they got. Perhaps most importantly, our analyses also reveal important differences between states: During the period we study, some of them experienced no policy change in the face of large conservative biases and significant changes in mass opinion (under-responsiveness); while others over-responded to changes in public opinion and as a result closed the representation gap. These patterns are likely to reflect a growing polarization in terms of policy outcomes, in the face of relatively similar trends in public opinion.

Needless to say, our approach still features a number of important limitations. First, like many studies in the existing literature (Lax and Phillips, 2009), our empirical analysis focuses on a single policy and as such its findings may not generalize to the opinion-policy relationship across all issues. Second, we are limited in the time-span of our data and as such we might underestimate responsiveness to the extent that the alignment of policies and preferences takes longer - as some of the mechanisms posited in the literature would require. Third, in our entire exercise we consider nominal minimum wage and as a result, some of our comparisons might understate or overstate differences across states or time periods. For instance, differences in nominal policy bias across states may be offset by corresponding differences in prices. That said, this distinction does not affect our claims on responsiveness.

In any case, our study could be taken as an example of how we can quantify the relationship between policies and mass preferences across more issues and a longer time-span in the United States and beyond. As our findings demonstrate, this could lead to important new insights on one of the most fundamental functions of democracies: translating individual preferences into laws.

References

- Ansolabehere, Stephen, and Brian F Schaffner. 2014. "Does survey mode still matter? Findings from a 2010 multi-mode comparison." *Political Analysis* 22 (3): 285–303.
- Caughey, Devin, and Christopher Warshaw. 2015. "Dynamic Estimation of Latent Opinion Using a Hierarchical Group-Level IRT Model." *Political Analysis* 23 (2): 197–211.
- Caughey, Devin, and Christopher Warshaw. 2018. "Policy Preferences and Policy Change: Dynamic Responsiveness in the American States, 1936–2014." *American Political Science Review* 112 (2): 249–266.
- Gelman, Andrew, Jeffrey Lax, Justin Phillips, Jonah Gabry, and Robert Trangucci. 2016. "Using multilevel regression and poststratification to estimate dynamic public opinion." *Unpublished manuscript, Columbia University*.
- Ghitza, Yair, and Andrew Gelman. 2013. "Deep interactions with MRP: Election turnout and voting patterns among small electoral subgroups." *American Journal of Political Science* 57 (3): 762–776.
- Gilens, Martin, and Benjamin I Page. 2014. "Testing theories of American politics: Elites, interest groups, and average citizens." *Perspectives on politics* 12 (3): 564–581.
- Lax, Jeffrey R, and Justin H Phillips. 2009. "Gay rights in the states: Public opinion and policy responsiveness." *American Political Science Review* 103 (3): 367–386.
- Lax, Jeffrey R, and Justin H Phillips. 2012. "The democratic deficit in the states." *American Journal of Political Science* 56 (1): 148–166.
- Ruggles, Steven, Sarah Flood, Sophia Foster, Ronald Goeken, Jose Pacas, Megan Schouweiler, and Matthew Sobek. 2020. "IPUMS USA: Version 10.0. Minneapolis, MN."
- Simonovits, Gabor, Andrew M Guess, and Jonathan Nagler. 2019. "Responsiveness without Representation: Evidence from Minimum Wage Laws in US States." *American Journal of Political Science* 63 (2): 401–410.

Online Appendix

Contents

A	Supplementary Information on Data	13
A.1	The seven public opinion surveys	13
A.2	Summary statistics of survey data	16
A.3	Additional data sources	18
B	Supplementary information on methods	19
B.1	Model specification	19
B.2	Estimating state-level minimum wage preferences in 2013	20
B.3	Multilevel regression coefficients	21
C	Supplementary analyses	22
C.1	Additional results	22
C.2	Validation of estimates	22
C.3	Accounting for the multimodal distribution of preferences	24

A Supplementary Information on Data

A.1 The seven public opinion surveys

Survey A was commissioned by Reason Foundation and the Arthur N. Rupe Foundation and was carried out by Princeton Survey Research Associates between December 4-8, 2013 to a sample of 1,011 American Adults. It used random digit dialing and interviewed respondents both using cellphones and landlines.

We construct our outcome of interest based on two questions in the survey (Q21B and Q22). First, respondents were asked: “Do you think the federal government should set a minimum wage, or not?” . Then, those answering in the affirmative were asked: “what do you think the federal minimum wage should be?”. The responses to the second question were top coded at \$50.

Survey B was commissioned by ABC News and the Washington Post and was carried out by Langer Research Associates and Capital Insight between December 12-15, 2013 to a sample of 1,005 American adults. It used random digit dialing and interviewed respondents both using cellphones and landlines.

We construct our outcome of interest based on the following question: “The minimum wage in this country is now seven dollars and 25 cents an hour. What do you think it should be?”. The responses were again top coded at \$50.

Survey C was an omnibus research survey shared generously by Erik Peterson. It was fielded to a sample of 2,645 American adults from an opt-in panel curated by Survey Sampling International (SSI) between March 6-11, 2014. SSI uses quota sampling to ensure that resulting samples resemble the US adult population on key marginals.

We construct our outcome of interest based on the following question: “The current federal minimum wage is \$7.25 an hour.If you had to choose, what amount would you prefer the minimum wage per hour to be?” Respondents were asked to enter a value in terms of dollars and cents in a text box.

Survey D was an omnibus research survey carried out by the SMaPP lab at New York University. It was fielded a national survey to 3,500 respondents from YouGov’s panel in the winter of 2016. To approximate a representative sample of the US adult population, YouGov uses matched sampling. This involves taking a stratified random sample of the

target population and then matching available Internet respondents to the target sample Ansolabehere and Schaffner (2014).

We construct our outcome of interest based on the following 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)? ” Respondents were asked to enter a value in terms of dollars and cents in a text box. Respondents’ state and their then current minimum wages were piped in based on their profile data.

In addition this survey also included a question measuring respondents’ preference about the federal minimum wage. That question read as follows: “The federal minimum wage is \$7.25 an hour. How much (in dollars) do you think it should be (0 meaning there should be none)?” We use responses to both questions.

Survey E was an omnibus research survey carried out by the authors of this study. It was fielded a national survey to 4,807 respondents sampled through the survey aggregator, Lucid Theorem in the winter of 2019. Lucid is a large marketplace for online survey panels, employing quota sampling to mirror the general population’s marginal distribution of age, gender, education, race and region.

We construct our outcome of interest based on the following question: “Currently, state minimum wages in the U.S. range from \$7.25 (the federal minimum wage) to \$12 (the highest state minimum wage in the country). In your view, what should be the minimum wage in your state? You can use 2 decimal places to specify cents (as in 7.25) and write 0 if you think there should be no state minimum wage.” Respondents were asked to enter a value in terms of dollars and cents in a text box.

We asked an analogous question about the federal minimum wage that read: “Currently, the federal minimum wage is \$7.25. In your view, what (in dollars) should be the federal minimum wage? You can use 2 decimal places to specify cents (as in 7.25) and write 0 if you think there should be no federal minimum wage.”

Surveys F and G were collected by YouGov applying quota sampling on gender, age, education and region. The two surveys combined reached 2,667 respondents. Just like in Survey E, respondents were asked: “Currently, state minimum wages in the U.S. range from \$7.25 (the federal minimum wage) to \$13.69 (the highest state minimum wage in the country). In your view, what should be the minimum wage in your state? You can

use 2 decimal places to specify cents (as in 7.25) and write 0 if you think there should be no state minimum wage.” Respondents then entered their minimum wage preference into an open field validated to include a number between 0 and 50.

Table OA1: Surveys and measures

Study	Survey firm	Year	N	Mode	Sampling	Federal m.w.	State m.w.
A	Reason/Rupe	2013	1,011	Phone	RDD	Yes	No
B	ABC	2013	1,005	Phone	RDD	Yes	No
C	SSI	2014	2,645	Online	Quota sample	Yes	No
D	YouGov	2016	3,500	Online	Matched sample	Yes	Yes
E	Lucid	2019	4,807	Online	Quota sample	Yes	Yes
F	YouGov	2021	1,350	Online	Quota sample	No	Yes
G	YouGov	2021	1,317	Online	Quota sample	No	Yes

A.2 Summary statistics of survey data

In addition to the items measuring preferences about the minimum wage - both state and federal in the case of the 2016 and 2019 data we use data on the following demographic variables:

- *Age*: Measured in years, then recoded to four age categories (18-29; 30-44; 44-64; 65+).
- *Gender*
- *Race*: Collapsed to a trichotomous indicator for White, Black and Hispanic with other responses coded as White
- *Education*: Recoded to four categories (no HS diploma; HS diploma; Some college; College degree; Post-grad degree)
- *State of residence*

Table OA2 below reports mean values for each of these demographic variables in each survey.

Table OA2: Summary statistics on demographics

	A: Reason	B: ABC	C: SSI	D: YouGov	E: Lucid	F:YouGov	G:Yougov	Total	ACS (2015)
18-29	14.2	10.2	16.3	11.1	19.9	24.4	23.5	16.5	21.62
30-44	24.1	16.8	24.6	28.3	27.8	24.7	25.2	27.1	25.25
44-64	36.5	38.7	38.2	40.3	33.1	33.6	33.5	36.2	33.88
65+	25.2	34.3	20.8	20.3	19.2	17.3	17.8	20.3	19.25
White+other	82.1	83.0	77.1	80.8	77.7	72.1	73.8	78.5	71.99
Black	9.6	9.4	11.9	11.4	11.3	14.0	13.6	11.5	12.96
Hispanic	8.3	7.6	11.0	7.8	11.1	13.9	12.6	10	15.05
No HS	4.4	5.3	2.3	2.9	4.0	6.9	5.7	3.7	12.9
High school	22.3	28.3	28.5	29.0	26.0	33.3	33.7	27.8	27.9
Some college	29.1	19.6	35.4	35.8	32.9	17.2	19.7	31.3	31.12
College	30.1	29.3	24.4	20.9	26.1	32.1	30.4	25.7	17.84
Postgrad	14.2	17.5	9.5	11.3	11.0	10.4	10.4	11.5	10.24
Male	49.7	51.7	49.1	44.2	49.3	52.4	48.1	47.9	48.63
Female	50.4	48.3	50.9	55.8	50.7	47.6	51.9	52.1	51.37

A.3 Additional data sources

- Data on minimum wage laws are obtained from
<https://www.dol.gov/agencies/whd/minimum-wage/state>
- Data on the Democratic party's two party vote share at presidential elections are obtained from
<https://electionlab.mit.edu/data>
- Data on rental prices are obtained from the Apartment List National Rent Report³
As the data goes back till January 2014, we used these earlier statistics for as the 2013 estimates

³<https://www.apartmentlist.com/rentonomics/national-rent-data/> (last accessed 2020-05-19).

B Supplementary information on methods

B.1 Model specification

Equations 1 and 2 describe the two multilevel regression models used to estimate individual level minimum wage preferences in Stage 1 of the MRP procedure. In both equations, respondents are indexed with the letter “i”. Our multilevel models relies on the default priors in the RStanArm software package (version 2.18.2). These weakly informative priors are adjusted to the data to help the model to converge without having an effect on the estimates.

Equation 1 estimates the probability that a respondent has a non-zero local minimum wage preference.

$$\begin{aligned} y_i &\sim \text{Bernoulli}(p_i) \\ \text{logit}(p_i) &= \alpha + \delta_1 \text{democratic_voteshare} \\ &\quad + \delta_{\text{state}[i]} \text{item} \\ &\quad + \beta_{\text{age}[i]} + \beta_{\text{sex}[i]} + \beta_{\text{edu}[i]} + \beta_{\text{race}[i]} \\ &\quad + \beta_{\text{state}[i]} + \beta_{\text{year}[i]} + \beta_{\text{state}[i], \text{year}[i]} \end{aligned} \tag{1}$$

Summary of prior distributions:

- Intercept (after predictors centered) $\sim \text{normal}(\text{location} = 0, \text{scale} = 2.5)$
- Coefficients (delta 1) $\sim \text{normal}(\text{location} = [0], \text{scale} = [27])$
- Covariance $\sim \text{decov}(\text{reg.} = 1, \text{conc.} = 1, \text{shape} = 1, \text{scale} = 1)$

Equation 2 estimates the preferred minimum wage among respondents with non-zero preference.

$$\begin{aligned}
y_i &\sim \text{Normal}(\mu_i, \sigma) \\
\mu_i &= \alpha + \delta_1 \textit{democratic_voteshare} \\
&\quad + \delta_{\textit{state}[i]} \textit{item} \\
&\quad + \beta_{\textit{age}[i]} + \beta_{\textit{sex}[i]} + \beta_{\textit{edu}[i]} + \beta_{\textit{race}[i]} \\
&\quad + \beta_{\textit{state}[i]} + \beta_{\textit{year}[i]} + \beta_{\textit{state}[i], \textit{year}[i]}
\end{aligned} \tag{2}$$

Summary of prior distributions:

- Intercept (after predictors centered) $\sim \text{normal}(\text{location} = 11, \text{scale} = 9.3)$
- Coefficients (delta 1) $\sim \text{normal}(\text{location} = [0], \text{scale} = [99])$
- Auxiliary (sigma) $\sim \text{exponential}(\text{rate} = 1/0.27)$
- Covariance $\sim \text{decov}(\text{reg.} = 1, \text{conc.} = 1, \text{shape} = 1, \text{scale} = 1)$

B.2 Estimating state-level minimum wage preferences in 2013

Note that our three surveys from 2013 lack measures of state-level minimum wage preferences, but include measures of federal minimum wage preferences. To get 2013 state preference estimates, we exploit the fact that in 2016 and 2019 we measure both state-level and federal preferences. These two latter surveys allow our model to estimate the difference between state and federal preference for each state. To get the 2013 state estimates, we simply make the assumption that this difference did not change between 2013 and 2016-19.

B.3 Multilevel regression coefficients

Figure OA1: Multilevel regression coefficients from the two models predicting minimum wage preferences

Note: Horizontal bars denote 67% - 89% - 97% credible intervals. Estimates for state and state-year effects are omitted from the plot

C Supplementary analyses

C.1 Additional results

Table OA3: Estimates of static responsiveness

	<i>Dependent variable:</i>			
		State Minimum Wage		
	2014	2016	2019	2021
Minimum Wage Preference	0.74 [0.52, 0.97] (0.53, 0.95)	0.97 [0.69, 1.24] (0.72, 1.22)	1.53 [1.11, 1.96] (1.15, 1.91)	2.26 [1.92, 2.61] (1.76, 2.76)
Constant	7.64 (7.52, 7.76)	8.02 (7.82, 8.21)	8.69 (8.37, 9.01)	9.19 (8.79, 9.59)
Observations	51	51	51	51
Adjusted R ²	0.48	0.53	0.55	0.61

Note: Uncertainty estimates denote 95% credible intervals [with] or (without) propagating uncertainty.

C.2 Validation of estimates

Figure OA2: Relationship between state level rent prices and minimum wage preferences.

Note: Each dot denotes a state, the blue lines are loess-curves.

How well our models can predict state level minimum wage preferences? As a validation exercise we verify that our estimates track changes in local rents, based on the simple assumption that if people notice that rents went up in their state in a given time period, they are also more likely to think that the minimum wage should be increased. Figure OA2 shows that in each year there is a close correspondence between rent and minimum wage preferences, with both rents and minimum wage preferences growing across the years. Yet, these cross-sectional relationships could be biased by several omitted variables.

As a more nuanced test, we regress our estimated preferences on local rent prices while including fixed effects (i.e. dummies) for state and year. Such a two-way fixed effects regression model zooms in on the within-state variance in rent over and above the national trends in rent prices. In other words, it controls for time-invariant confounders (omitted variables whether observed or unobserved, that do not change within a state in our study period) and state-invariant confounders (omitted variables which affect each state in a given year in the same way). Following best practices in the literature, we scale rent prices to a large but realistic (95th percentile) within state deviation from national trends at \$80.⁴ The regression estimates in Table OA4 indicate that if rent prices increased by \$80 in a state, hourly mean wage preferences also increased by 11 cents. Put differently, taking an average of ($40\text{hours} \times 52\text{weeks}/12\text{months} =$) 173 working hours per month, a full time minimum wage worker's monthly salary is preferred to increase by \$19 for a \$80 increase in rents.

Table OA4: 2FE regression of minimum wage preferences on rent prices across US states

	<i>Dependent variable:</i>
	Mean MW preference
Rent	0.11** (0.04)
Observations	204
Adjusted R ²	0.94
Residual Std. Error	0.28 (df = 149)
<i>Note:</i>	*p<0.05; **p<0.01; ***p<0.001

⁴Mummolo, Jonathan, and Erik Peterson. "Improving the interpretation of fixed effects regression results." Political Science Research and Methods 6.4 (2018): 829-835.

C.3 Accounting for the multimodal distribution of preferences

In this section we outline the importance of accounting for the multimodal distribution of minimum wage preferences. Recall that our surveys constrain the answers to a range of \$0-\$50. The top panel of Figure OA3 plots the distribution of these answers. It is apparent that a large majority of respondents reveal a preference between \$7 and \$20. Within this range, we see spikes in the distribution at \$10, \$12 and \$15, which stems from respondents' well-known tendency provide round numbers. Much more consequential for our analysis is the spike at \$0 that shows that almost 10% of respondents think their state ought not to have a minimum wage law.

The second and third panels of Figure OA3 show 50 draws from the posterior predictive distributions from two multilevel regression models. The second panel is from a naive linear model that ignores the clustering around \$0. Meanwhile, the third panel shows the posterior combining two models, one which estimates the likelihood that a respondent reveals a non-zero preference, and another which estimates the preference among those with non-zero preference. Comparing panels two and three demonstrates that the latter model offers a much better fit. The naive regression underestimates the share of respondents with a \$0 preference, but compensates with a rather wide normal distribution, which predicts way too many respondents to have an ideal minimum wage policy below \$0 or above \$20. Conversely, the model we apply in our analyses not only accounts for people with a \$0 preference, but mirrors the non-zero preferences much more closely. Notice too that compared to the third panel, the center of the distribution in the second panel is pulled closer to 0.

Figure OA3: Posterior predictive distributions

Figure OA4: State-level minimum wage preference estimates based on two multilevel models

Figure OA5: Juxtaposing mean predictions and their uncertainties from the two models

How do these modeling decisions affect our estimates of state-level minimum wage preferences? We juxtapose the poststratified estimates for each state in each year we study in Figure OA4. Two trends are imminent. First, the naive model (blue dots) predicts consistently lower minimum wage preferences than the model accounting for the bimodal distribution (red dots). This conservative bias amounts to 80 cents on average. Second, the naive model is also overly confident in its predictions. Figure OA5 offers an alternative visualization of the comparison between the two model predictions with each state-year denoted with a single black point and a linear regression line in blue. Panel A demonstrates that the negative bias in the mean predictions is consistently around 80 cents. This is a substantively important difference considering that one of our main substantive finding is that the average conservative policy bias was between \$2.2 and \$3.1. Had we used the naive model's predictions, we would have concluded that the policy bias was between 26-36% lower than it is. Perhaps even more concerning, Panel B shows that the naive model is overly confident of its estimates. The average width of a 90% credible interval is 30% smaller in the naive model than in the bimodal model, although for some state-years the uncertainty is three times larger once we account for clustering around \$0.