The illusion of responsiveness: Evidence from minimum wage laws*

Gabor Simonovits† Central European University

Alexander Bor‡ Aarhus University

July 29, 2020

Abstract

To what extent does policy respond to changes in mass policy preferences in U.S. states? Current approaches fail to quantify dynamic (i.e. temporal) responsiveness because they lack meaningful, comparable scales or repeated measures on attitudes and policy. We overcome this issue by measuring minimum wage preferences and policies on the same cardinal scale combining novel and secondary public opinion data (N = 17,619). Relying on multilevel regression and poststratification, we estimate Americans' minimum wage preferences and compare them to observed policies in each state from 2013 to 2019. Fixed-effects regressions demonstrate that, on average, policy very closely tracked changes in preferences, but the size of policy bias remained relatively stable. Minimum wage laws were consistently \$1.75–\$2.25 below citizens' mean preference.

*We thank Andy Guess for his work on initial phases of his project. We are also grateful for comments from seminar participants from 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

The illusion of responsiveness:
Evidence from minimum wage laws

Are policies responsive to changes in mass preferences? Most normative accounts of democratic politics argue that they should be (Dahl, 1973), and some classic analytical models imply that they are (Downs, 1957). At the same time, empirical research on contemporary democratic politics casts serious doubts on the veracity of these theories. First, given the evidence of massive divergence across the platforms of major parties (Bafumi and Herron, 2010), one might argue that even small changes in electoral preferences could lead to drastic policy change due to the impact of changing party control (Caughey, Xu, and Warshaw, 2017). In this case, we would expect policy outcomes to be overly responsive to public opinion. Another perspective holds that due to the outsized influence of the wealthy and the power of interest groups, changes in the preferences of ordinary citizens are irrelevant when it comes to policy making (Gilens and Page, 2014). This view is also supported by studies showing that state policies are as often than not are supported by opinion majorities (Lax and Phillips, 2012).

Against the backdrop of this debate, it seems essential to *quantify* the association between opinion change at the mass level and changes in policy outcomes. Studies of policy responsiveness fall short of achieving this, mostly due to a lack of appropriate data and research designs. For instance, while some studies quantify the cross-sectional association between public opinion and policy outcomes (Lax and Phillips, 2012), they lack repeated measures to explore whether and to what extent *changes* in preferences matter for policy. Meanwhile, studies that do track changes in opinion and policy outcomes at the national (Wlezien, 1995) or the subnational level (Caughey and Warshaw, 2018) use different, non-comparable scales. Thus, while there is some evidence that changes in mass opinion are associated with *some* change in policy outcomes, there is no evidence on the *magnitude* of this association.

The reason it is important to quantify the *magnitude* of dynamic responsiveness – beyond simply rejecting the null hypothesis of no responsiveness – is that it is necessary to distinguish between scenarios with starkly different normative appeal. First, under-responsiveness – i.e. policy change that is slower than opinion change – would lead policy outcomes to remain far removed from the average preferences insofar as the status quo falls far from what most people prefer. Second, hyper-responsiveness – i.e. policy change that is faster than opinion change – could lead policy outcomes to leapfrog around the preferences of relatively moderate electorates. Crucially, current approaches to the study of dynamic representation fail to distinguish between these possibilities empirically.

In this paper we deploy a research design and marshal data that is appropriate to quantify the dynamic relationship between mass opinion and policy – at least in one specific issue domain. Like Simonovits, Guess, and Nagler (2019), we focus on minimum wage laws, a context that has been widely used to study both legislative politics and representation (Clinton, 2012). Our analysis relies on five national surveys that each measure preferences about the minimum wage using open-ended questions (i.e. preferences are measured in U.S. dollars). Using an augmented dynamic multilevel model with poststratification (MRP), we obtain estimates of average preferences for each U.S. state in the years 2013, 2016 and 2019 and match the resulting data with minimum wage laws in the same years. We use the resulting panel data set to make inferences on the opinion-policy relationship using a difference in differences design.

Consistent with existing research (Caughey and Warshaw, 2018), we find evidence suggesting that policy-change and preference change are indeed associated. However, we also find that state policy – at least in the case of minimum wage laws – is *under-responsive.* According to our models, a dollar increase in average preferences regarding the minimum wage results in a 10-25-cent increase in state policy – depending on the pairs of years studied. Moreover, our estimates are precise enough to rule out a magnitude of dynamic responsiveness exceeding 50 cents even in our most optimistic specifications. Using *policy bias* as an alternative measure of representation (Simonovits, Guess, and Nagler, 2019), we show that under-responsiveness results in policy outcomes that remain far removed from corresponding preferences with policy gaps between –2.5 and +2 dollars.

These findings have important implications for the study of representation and policy change. Most importantly, they again highlight the importance of careful measurement for making normative claims about the nature of respresentation. In this particular case, our results indicate that while policy *does* respond to changes in mass opinion, the extent of responsiveness is too small to align outcomes with preferences. Of course, it is entirely possible that our results are specific to the particular policy environment that we study. However, proving that would require more effort to *quantify* policy representation in other domains, too.

Research design

Our empirical analysis explores the relationship between changes in average preferences about the minimum wage and corresponding changes in state laws between 2013 and 2019. Our focus on a single policy issue – rather than a composite measure of policy outcomes – 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.

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 12 statewide ballot initiatives have been voted on. In short, we minimum wage policies constitute a most likely case for finding substantial dynamic responsiveness. Finally, quantifying representation on this issue is easy as changes in policy outcomes and preferences can be measured on a meaningful and comparable cardinal scale.

Data

We rely on five national surveys that share the unusual feature that respondents were asked to express their preferred levels of the minimum wage in dollars. We summarize information about the three datasets in Table 1 and provide additional details about them in the Online Appendix including (1) question wording, (2) representativeness and (3) minor coding decisions. Crucially, in the "supersurvey" we created from these data, we have 19,452 responses to an item that asks respondents what the minimum wage should be. Additionally, we use the following covariates in our analyses: state of residence, gender, race/ethnicity, age, education and partisan identification.

Table 1: 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	1005	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

Note that the five surveys diverge in terms of size, mode and sampling. In particular, three of the five surveys rely on online opt-in panels. However, because we rely on multilevel modeling and poststratification, our state-level estimates of public opinion could well approximate true state-level preferences if inclusion to these samples were ignorable to our dependent variable – i.e. minimum wage preferences – conditional on the demographic variables we include in our models. In the Online Appendix, we present ample evidence of the validity of our state-level estimates. In each year, we observe a strong correlation between our estimates and rental prices as well as general economic ideology estimates. Moreover, our estimates of *preference changes* are correlated with changes in the aforementioned variables.

Estimating state-level preferences

To obtain state-specific preferences for each of the three years we study, we follow recent literature (Ghitza and Gelman, 2013; Lax and Phillips, 2012; Kastellec et al., 2015) 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.

In stage 1 of the MrP, we build a multilevel model predicting each respondent's preferred level of minimum wage. Our model includes standard demographic variables (age, sex, education, race and respondent state) and state-level predictors including Obama's 2012 two-party vote share, state-level economic ideology provided generously by Caughey and Warshaw (2018) as well as average rental prices (for more details on our data and estimates, see Sections A & C in the OA, respectively).

To account for the dynamics of public opinion, we include both year and state-year variables to account for time effects. Our approach is superior to running separate models for each year because it allows partial pooling across time. This is crucial for estimating state-level minimum wage preferences from surveys with only federal minimum wage preferences.[1] Moreover, economic ideology estimates (Caughey and Warshaw, 2018) and average rental prices vary across both states and years, and thus we can use that to pool information about time trends.

Finally, we add a random intercept representing responses to questions on the *state* versus the *federal* minimum wage. Technically, the multilevel model predicts responses to both survey items, and in the two online surveys we use both survey responses of each survey taker. In other words, we use surveys including both versions of the question to estimate wording effects and extrapolate to preferences about state minimum wages in all surveys. We observe in total 19,458 responses from 12,968 survey takers.

We implement our MrP analysis via Rstanarm, an R library that uses Bayesian simulation to estimate complex multilevel models. The simulations result in posterior draws for each of our parameters and thus we can also obtain posterior distributions to our MrP estimates of average preferences in each state and year pair. Our main model is specified as follows[2]:

$$
\begin{aligned}
y_i &\sim Normal(\mu_i, \sigma) \\
\mu_i &= \alpha + \delta_1 obama\_2012 + \delta_2 state\_rent + \delta_3 ideology \\
&\quad + \delta_{state[i]} item \\
&\quad + \beta_{age[i]} + \beta_{sex[i]} + \beta_{edu[i]} + \beta_{race[i]} \\
&\quad + \beta_{state[i]} + \beta_{year[i]} + \beta_{state[i],year[i]}
\end{aligned}
\tag{1}
$$

Once we build our multilevel model, we can simulate the state-level minimum wage preference for all possible combinations of our independent variables (demographic groups, states and years). In essence, this yields an estimate for minimum wage preference from all 55,080 types of US citizens our data can capture. Finally, these predictions are weighted (post-stratified) to the US Census to account for the unique demographic composition of each state in each year and aggregated to state years.

[1] While studying state minimum wage preferences spanning six years is a major step forward in quantifying dynamic responsiveness, having only three points in time is insufficient to incorporate time smoothers, as recommended by (Caughey and Warshaw, 2018).

[2] For our priors see Section B.1 in the Online appendix

Inference of dynamic responsiveness

Our goal is to obtain estimates of dynamic responsiveness, i.e. the association between preference change and policy change. 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 2013, 2016 and 2019. Based on the resulting data, we estimate linear regressions predicting state minimum wages with average preferences in the corresponding state-year observations and fixed effects for states and years (Caughey and Warshaw, 2018). While our difference-in-differences approach rules out time-invariant confounders that could lead to a cross-sectional correlation between preferences and policy even in the absence of dynamic responsiveness, our estimates should not be interpreted as causal.

A final key element of our empirical approach is to account for the measurement error in our independent variable. Given that we do not directly observe average policy preferences but rather infer them from an MrP model, measurement error in these unobserved quantities could bias point estimates and standard errors.[3] Therefore, we account for measurement error using an approach known as the "method of composition" (MOC) or "propagated uncertainty" (Caughey and Warshaw, 2018; Kastellec et al., 2015). Like other applications, these adjustments attenuate the estimated effects of preference on policy quite significantly.

Results

We start by considering static (i.e. cross-sectional) responsiveness in each year included in our study. The three panels in Figure 1 plot state policies in 2013, 2016 and 2019 against average preferences in the corresponding states and years. Three clear patterns stand out. First, in each of the three years, policy outcomes were strongly related to variation in preferences; states with higher preferences had on average higher policies. Second, we observe variation across years: both preferences (medians: 2013–\$9.0; 2016–\$10.1; 2019–\$10.2) and policies (medians: 2013–\$7.2; 2016–\$8.1; 2019–\$8.8) increased between 2013 and 2019. Figure OA5 in the OA offers details on the state-by-state estimates of minimum wage preferences and policies, providing further evidence that the changes are

[3]Measurement error could arise both because our weighting procedure does not completely "unbias" our estimates of state preferences or because our estimates of cross-sectional and temporal variation in opinion are attenuated due to partial pooling.

quite substantial especially between 2013 and 2016. Finally, in each year, we observe a substantive conservative policy bias; almost every point in our figure is below the diagonal.

Figure 1: Static policy responsiveness by year

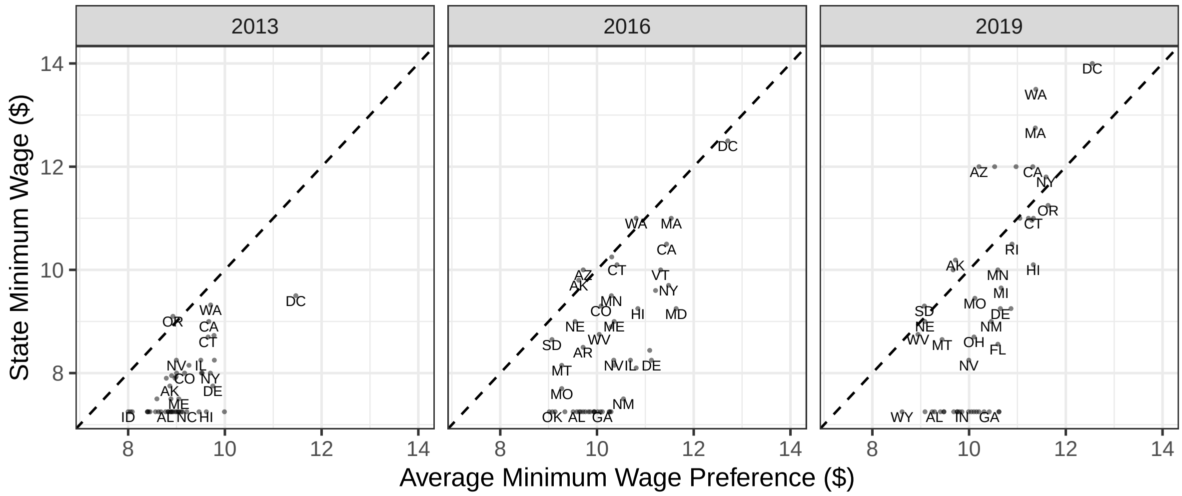

Figure OA6 offers more detail on changes in policy bias. It shows that the median bias grew from $1.4 in 2013 to $2.1 in 2016 and then returned to $1.4 in 2019. Meanwhile, policy bias in minimum wage laws also reflects an increasing polarization in the US. The range of biases grew from $2.9 to $3.3 in 2016 and $5.5 in 2019. Oregon had the only liberal bias in 2013 at 17 cents. Three states (Alaska, Washington and Arizona) had liberal biases in 2016, all below 30 cents. However, eleven states had liberal biases by 2019 with the highest bias in Washington above $2.1. Note that we make these comparisons in nominal terms. Adjusting for inflation – amounting to 9.7% between 2013 and 2019 - would leave our substantive conclusions unchanged.

Next, we report our main results on dynamic responsiveness in Table OA2. We quantify dynamic responsiveness by estimating linear regressions of policy change on preference change with state and time fixed effects. As explained above, we report models that propagate the uncertainty due to the measurement error in our explanatory variable (i.e. minimum wage preferences). The results show that across all years, a dollar increase in average preferences was associated with a 13-cent increase in state minimum wages. We find that dynamic responsiveness is smaller for the two three-year periods (2013-16, 2016-19) at 8 to 10 cents' increase in policy for every dollar change in preferences. Meanwhile, our estimate is larger for the six-year period between 2013 and 2019 at 24 cents predicted change. This makes intuitive sense as a longer timespan allows larger preference changes to accumulate, making it more likely that the legislators notice and have time to respond to them. We plot median preference change against minimum wage change for each state in each time period in Figure OA7.

Table 2: Estimates of dynamic responsiveness

	All years	2013-2016	2016-2019	2013-2019
Point estimate	0.13	0.1	0.08	0.24
80% CI	[-0.03, 0.32]	[-0.10, 0.3]	[-0.11, 0.26]	[-0.08, 0.55]

Our findings indicate that even if policy outcomes were perfectly aligned with average preferences in each state at one point in time, policy change would be too slow to track changes in public opinion. Note, however, that policies were already biased in 2013, perhaps due to under-responsiveness. Another way to assess dynamic responsiveness is to compare actual policy change with the counterfactual policy change that would have been necessary to align policy changes to preferences. In the final analysis in Figure 2, we plot changes in minimum wage as a function of *ideal change*, defined as the policy change required to perfectly track public preferences, without any bias.[4]

Figure 2: Policy change as a function of a normatively ideal change

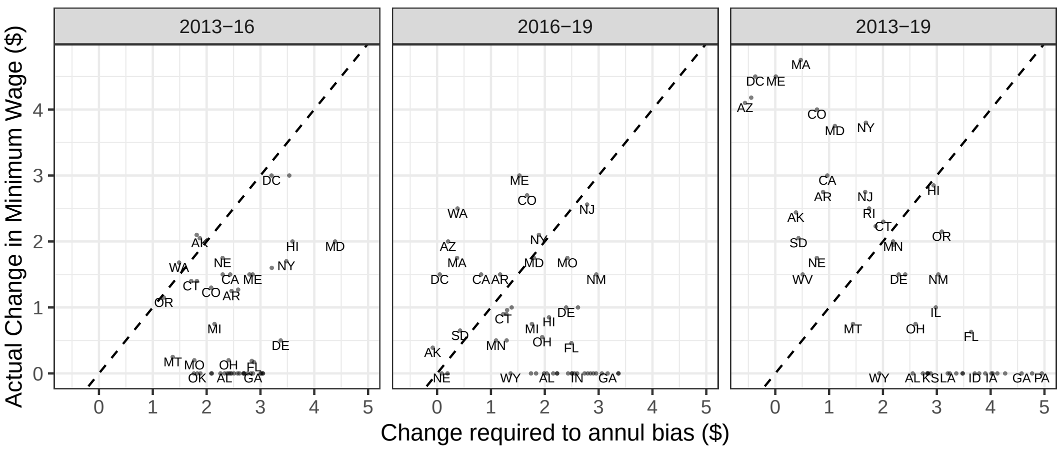

These analyses paint a grim picture of the state of dynamic policy responsiveness. First, they show that in the majority of states, policy change was too small compared to what was needed to eliminate policy bias (remember that several states had no changes in minimum wage). Second, we find important heterogeneity across the two periods we

[4]Formally, we define this quantity as the sum of policy bias in t1 and change in preferences between t1 and t2, multiplied by -1.

study. From 2013 to 2016, in all states except Alaska, Arizona and Washington, policy change was too conservative with the median -$2.1. In the second part of the time period, eleven states enacted policy change that exceeded what was necessary to eliminate bias. Indeed, the policy changes between 2013 and 2019 reveal a surprising pattern: States with smaller initial biases were more likely to increase their minimum wage laws, while states with large baseline conservative biases continued to rely on the federal minimum wage, which has not changed since 2009. Overall, these findings are consistent with the account that while policy change is indeed related to changes in mass preferences, dynamic responsiveness is too weak to align the two (see also Figure OA8 demonstrating the absolute level of policy bias).

Discussion

Our study offers new insights on the state of dynamic policy responsiveness 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 policy outcomes on the same scale allows us to investigate how such responsiveness translates to policy bias. In other words, we can explore whether the co-movement of opinion and policy can reduce the wedge between the two.

Based on our analysis of minimum wage laws, it appears that statehouse democracy is "responsive" in the sense that changes in policy outcomes are related to fluctuations of mass opinion within states. At the same time, we find that on average state policy is under-responsive: a dollar increase in the average preference for minimum wages translates to about a 13-cent increase in actual minimum wage laws. Moreover, we find substantial and persistent divergence between average preferences and policy outcomes in each year we study.

Our analyses also hint at heterogeneous dynamic processes, with some states experiencing no policy change in the face of large conservative biases and significant changes in mass opinion (under-responsiveness); while others implementing overly aggressive policy change compared the level of bias and moderate movements in public opinion. In this particular case, however, under-responsiveness and conservative bias appears to be the dominant pattern.

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 polices 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 public 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.

Bafumi, Joseph, and Michael C Herron. 2010. “Leapfrog representation and extremism: A study of American voters and their members in Congress.” *American Political Science Review* 104 (3): 519–542.

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.

Caughey, Devin, Yiqing Xu, and Christopher Warshaw. 2017. “Incremental democracy: The policy effects of partisan control of state government.” *The Journal of Politics* 79 (4): 1342–1358.

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

Dahl, Robert Alan. 1973. *Polyarchy: Participation and opposition.* Yale University Press.

Downs, Anthony. 1957. "An economic theory of political action in a democracy." *Journal of Political Economy* 65 (2): 135–150.

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.

Kastellec, Jonathan P, Jeffrey R Lax, Michael Malecki, and Justin H Phillips. 2015. "Polarizing the electoral connection: partisan representation in Supreme Court confirmation politics." *The journal of politics* 77 (3): 787–804.

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.

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.

Wlezien, Christopher. 1995. "The public as thermostat: Dynamics of preferences for spending." *American journal of political science* 39: 981–1000.

Online Appendix

Contents

A Supplementary Information on Data

A.1 The five 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: "hat 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."

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 thricotomous 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)
- *Partisan identification*: Democrats, Republicans, Independents (with leaners included with their respective parties)
- *State of residence*

Table OA1 below reports mean values for each of these demographic variables in each survey.

A.3 Additional data sources

- Data on minimum wage laws are obtained from `https://www.dol.gov/agencies/whd/minimum-wage/state`
- Data on Obama's 2012 two party vote share are obtained from `https://electionlab.mit.edu/data`
- Data on state level economic ideology are obtained from Caughey and Warshaw (2018).
- Data on rental prices are obtained from the Apartment List National Rent Report[5] As the data goes back till January 2014, we used these earlier statistics for as the 2013 estimates

[5] `https://www.apartmentlist.com/rentonomics/national-rent-data/` (last accessed 2020-05-19).

Table OA1: Summary statistics on demographics

	A: Reason	B: ABC	C: SSI	D: YouGov	E: Lucid	Total	ACS (2015)
18-29	14.2	10.2	16.3	11.1	19.9	15.6	21.62
30-44	24.1	16.8	24.6	28.3	27.8	26.2	25.25
44-64	36.5	38.7	38.2	40.3	33.1	36.8	33.88
65+	25.2	34.3	20.8	20.3	19.2	21.4	19.25
White+other	82.1	83.0	77.1	80.8	77.7	79.2	71.99
Black	9.6	9.4	11.9	11.4	11.3	11.1	12.96
Hispanic	8.3	7.6	11.0	7.8	11.1	9.6	15.05
No HS	4.4	5.3	2.3	2.9	4.0	3.5	12.9
High school	22.3	28.3	28.5	29.0	26.0	27.2	27.9
Some college	29.1	19.6	35.4	35.8	32.9	32.8	31.12
College	30.1	29.3	24.4	20.9	26.1	24.9	17.84
Postgrad	14.2	17.5	9.5	11.3	11.0	11.6	10.24
Male	49.7	51.7	49.1	44.2	49.3	48.1	48.63
Female	50.4	48.3	50.9	55.8	50.7	52.0	51.37
Republican	38.6	48.5	36.4	32.7	36.7	36.6	-
Independent	14.5	7.3	14.5	15.6	17.2	15.3	-
Democrat	46.9	44.2	49.1	51.8	46.1	48.1	-

B Supplementary information on methods

B.1 Weakly informative priors in Bayesian MRP

Our mutlilevel model 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.

- Intercept (after predictors centered) $\sim$ normal(location = 0, scale = 45.94)
- Coefficients (delta 1-3) $\sim$ normal(location = [0,0,0], scale = [1.32,5.27,8.84])
- Auxiliary (sigma) $\sim$ exponential(rate = 1/4.59)
- Covariance $\sim$ decov(reg. = 1, conc. = 1, shape = 1, scale = 1)

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

WE DON'T KNOW THIS, BUT WE KNOW FEDERAL PREFERENCES IN 2013 AND WE KNOW DIFFERENCE BETWEEN FEDERAL AND STATE PREFERENCE 16 AND 19, SO WE ASSUME THIS DIFFERENCE IS THE SAME IN 2013.

C Supplementary analyses

C.1 Validation of state-level predictors

Our MRP models rely on state level smoothers, most notably an estimate of state level economic liberalism and average rent. Best practice in the MRP literature show that these state level predictors improve estimation by increasing the level of partial pooling in the model. Below we validate our choice of these predictors by plotting the liberalism scores and rent prices against the predicted state level minimum wage preferences for each year (Figures OA1 and OA2), as well as the change in these indicators against the change of minimum wage preferences for each period in our data (Figures OA3 and OA4). These plots indicate a strong association between minimum wage preferences and state level economic liberalism and rent, respectively.

Figure OA1: Relationship between state level economic ideology and minimum wage preferences

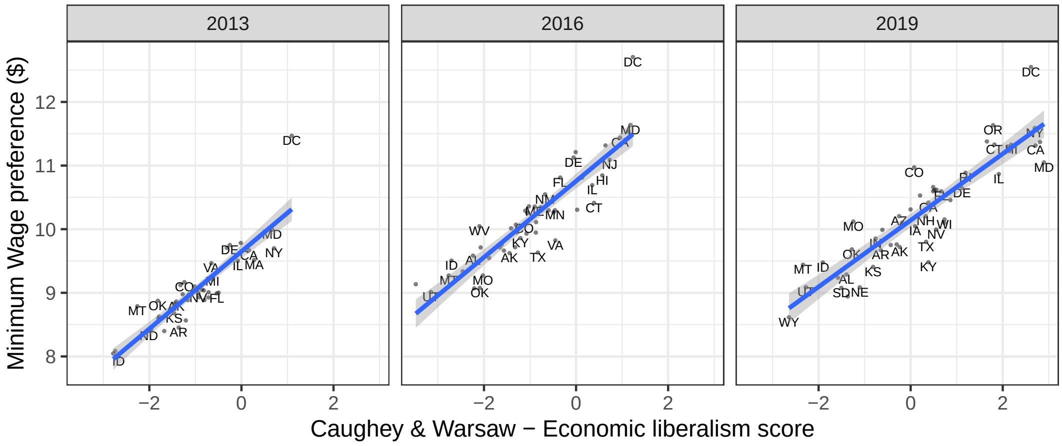

Figure OA2: Relationship between state level rent prices and minimum wage preferences

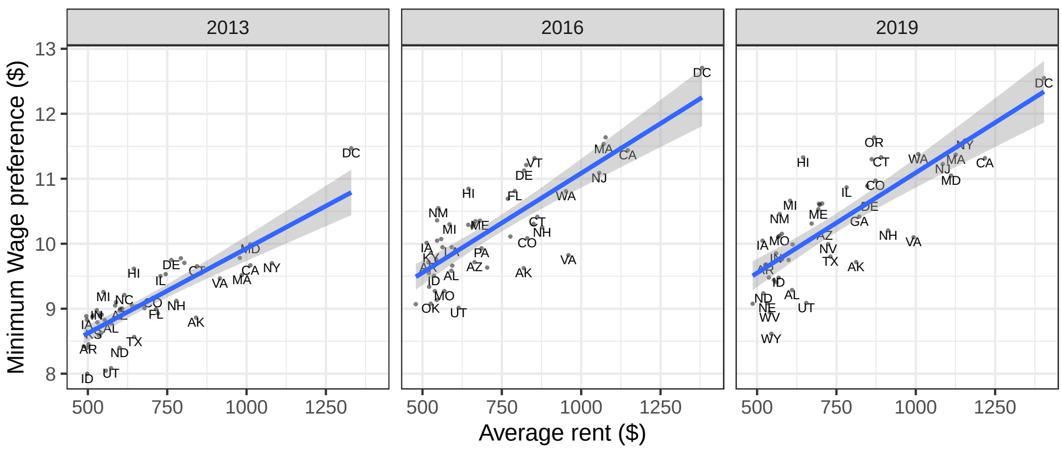

Figure OA3: Relationship between state level economic ideology change and change in minimum wage preferences

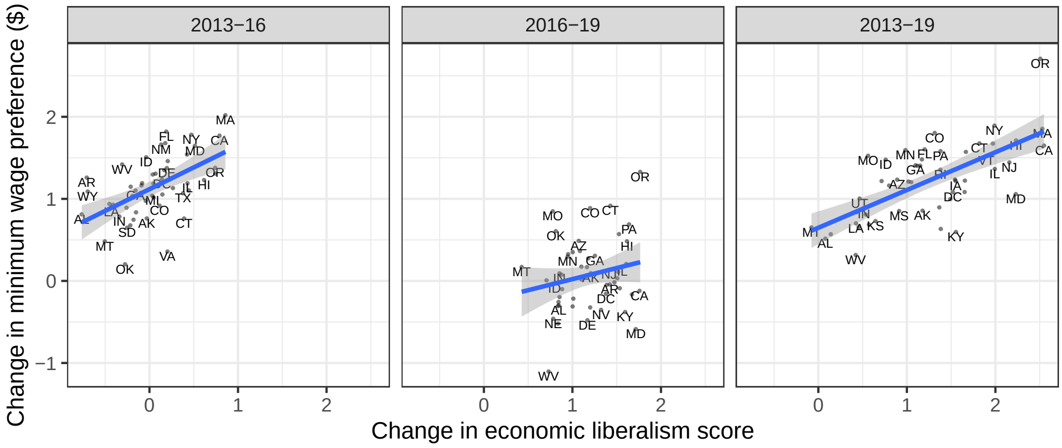

Figure OA4: Relationship between changes in state level rent prices and minimum wage preference change

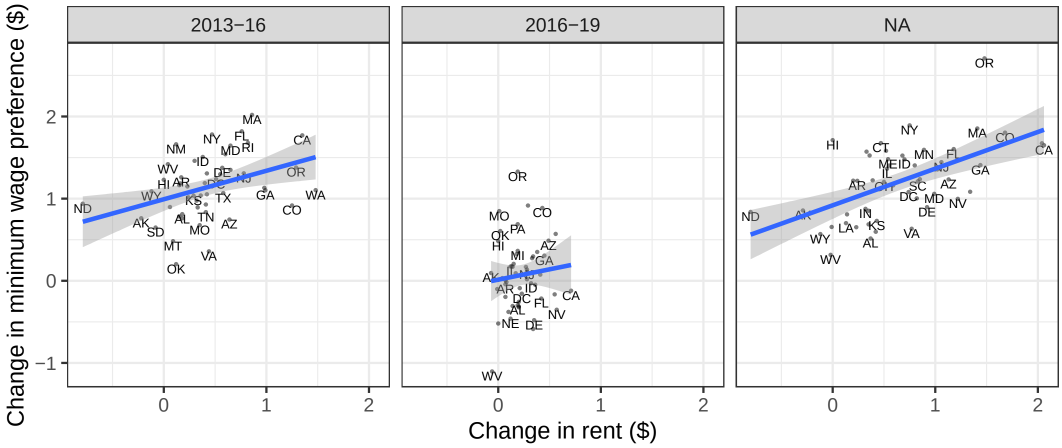

C.2 Policy and preference change

Figure OA5 offers a more detailed look at policy change (in red) and predicted preference change for each year (x) and state in our data.

Figure OA5: State level minimum wage estimates (black) and policies (red) by year

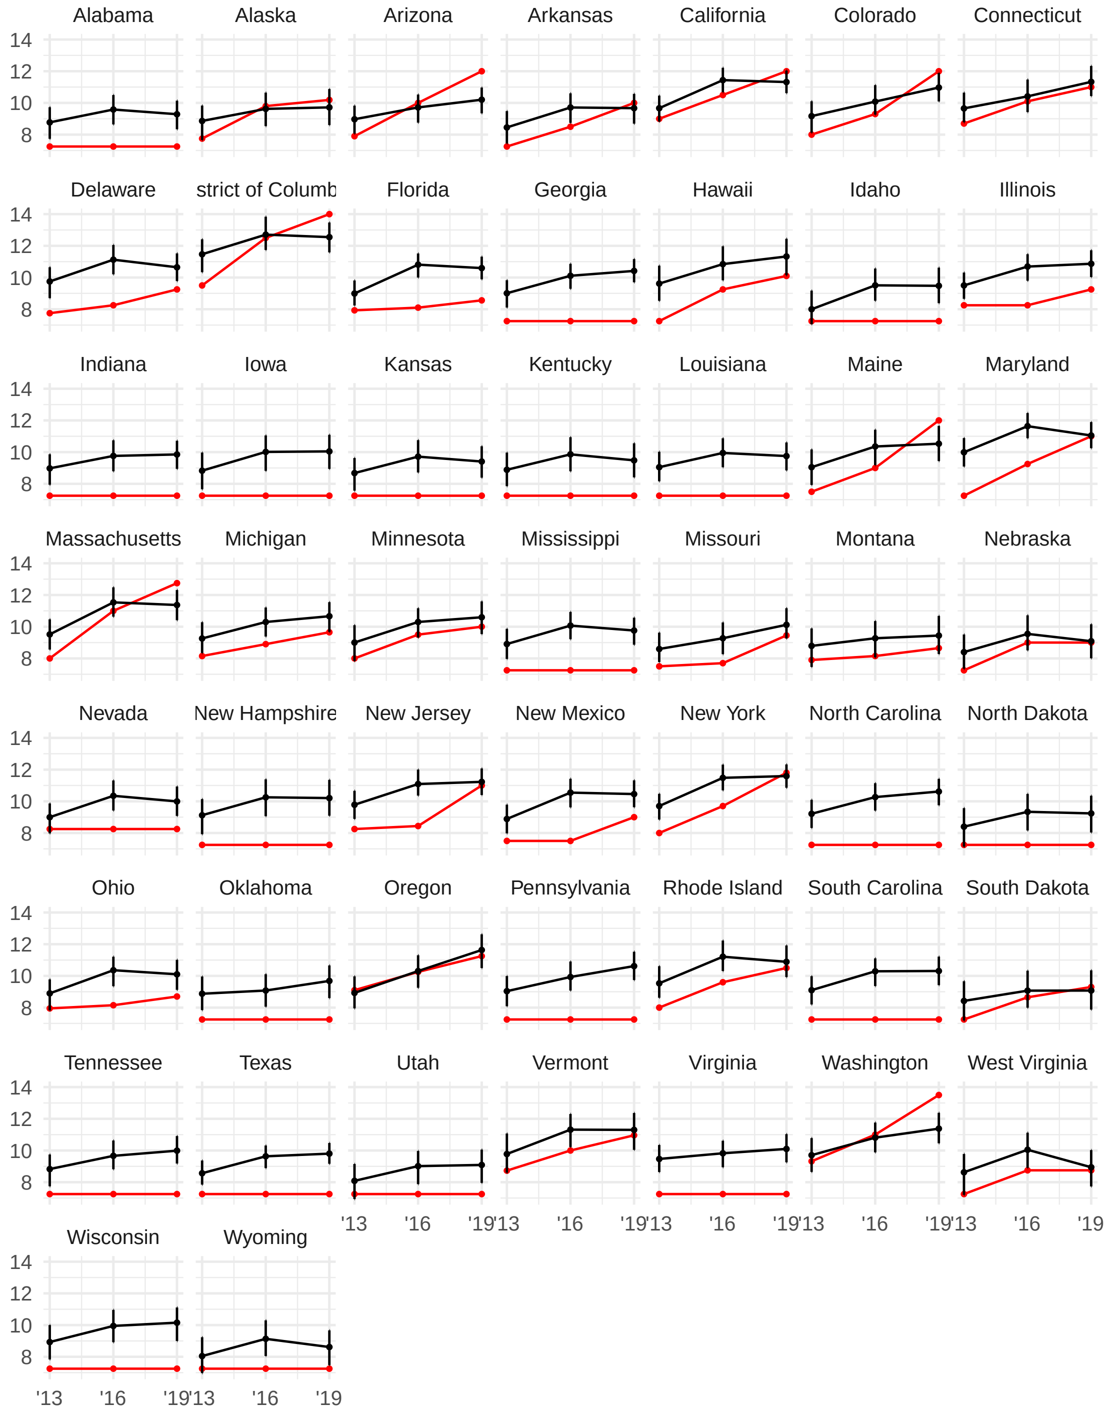

Note: Errorbars denote 95% credible intervals

C.3 Static policy bias

Figure OA6 offers an alternative view of policy bias across the three years in our data by displaying the distributions. This plots highlight two important patterns: 1) median policy bias was highest in 2016 and about the same in 2013 and 2019. 2) the range of policy biases have increased considerably in every year in our data.

Figure OA6: The distribution of state-level policy bias (policy - preference) by year

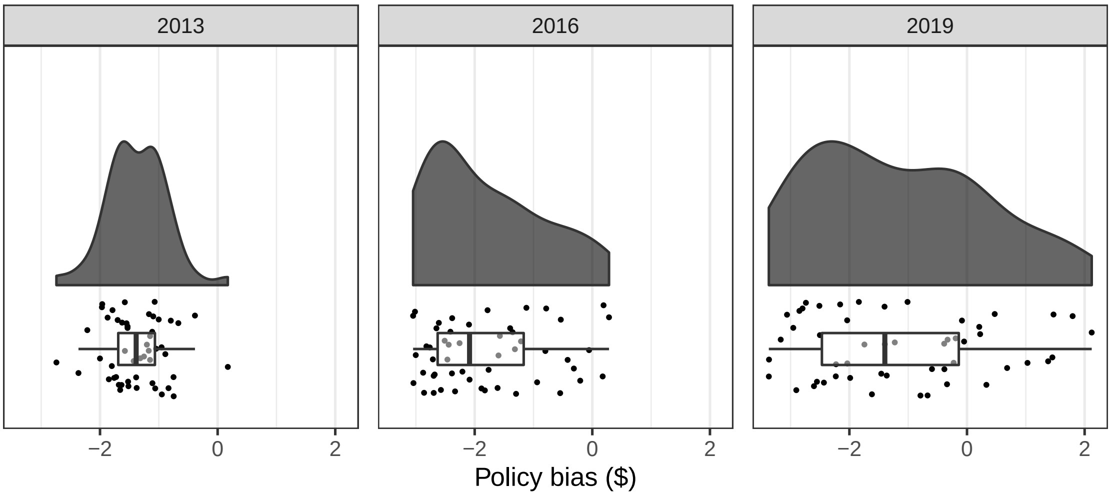

C.4 Plot dynamic responsiveness

Figure OA7: Dynamic responsiveness: change in median preferences (x) and in policies (y)

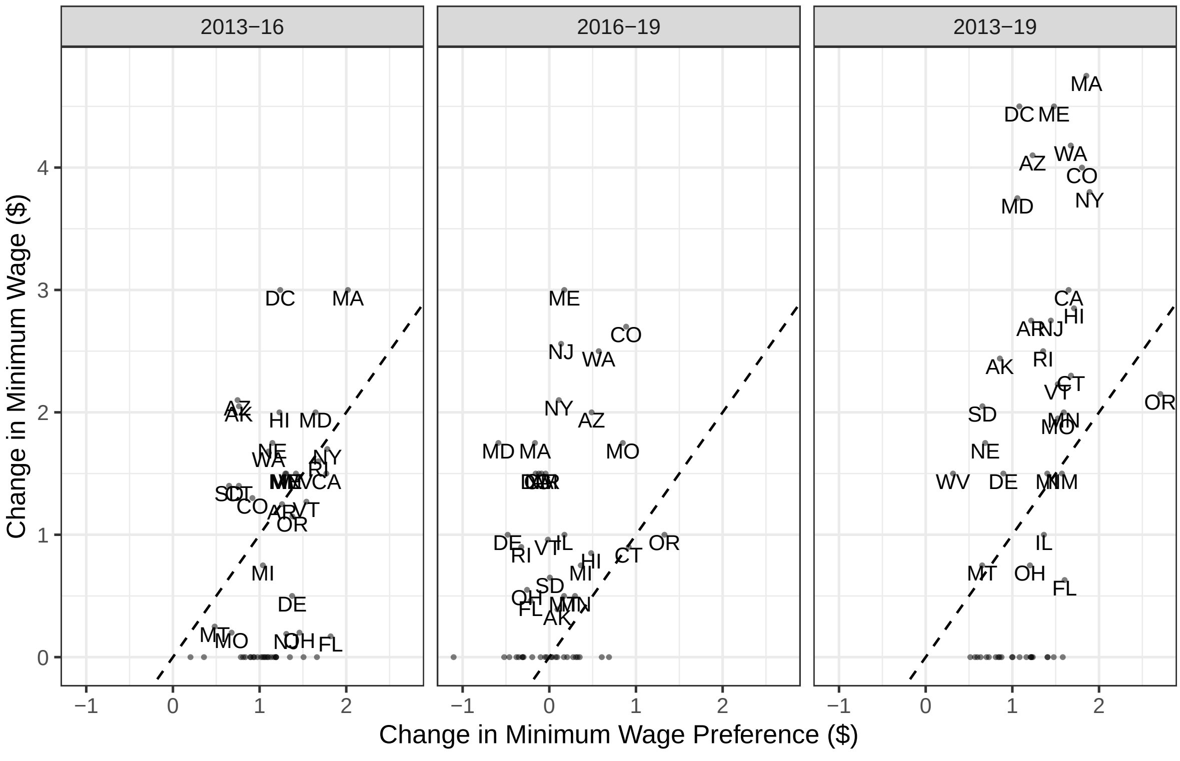

C.5 Change in absolute policy bias

Figure OA8 offers an alternative look on the state of dynamic responsiveness by displaying changes in absolute levels of policy bias. From a normative perspective, we would hope that the level of a absolute policy bias decreases year after year. Yet, the plot clearly shows that this is not the case and increases in bias (whether in conservative or liberal direction) are more frequent than decreases.

Figure OA8: Changes in the absolute level of policy bias for each period

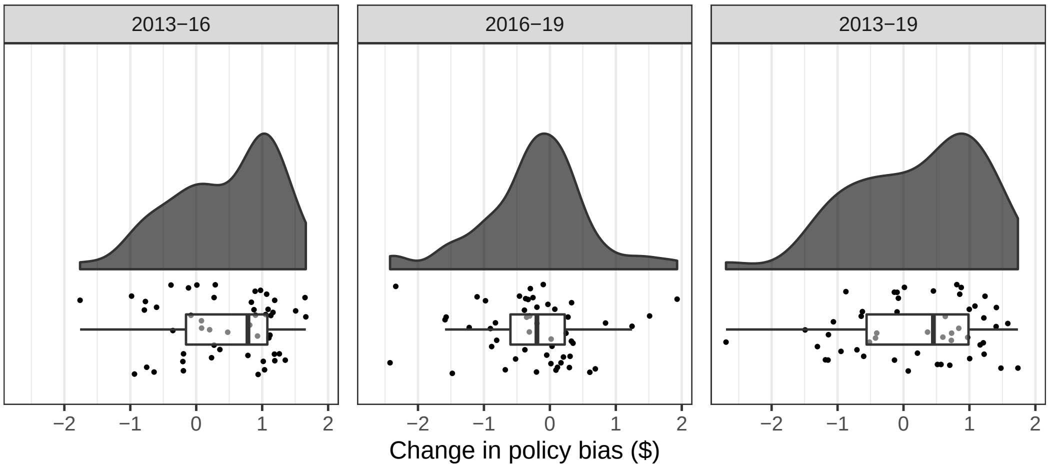

C.6 Quantifying static responsiveness

Table OA2: Estimates of static responsiveness

	2013	2016	2019
Point estimate	0.24	0.58	0.95
80% CI	[0.09, 0.38]	[0.33, 0.86]	[0.63, 1.29]