Online Supplementary Material for "Beyond the Mean: How Thinking About The Distribution of Public Opinions Reduces Politicians' Perceptual Errors"

December 1, 2025

Contents

1 Survey Data (General Public)

To optimise the precision and accuracy of our estimates of municipal ideology, we collected public opinion data from as many recent Canadian surveys as possible. We describe these data sources in Table SM.1. Our main source of survey data is the Consortium on Electoral Democracy (C-DEM), a major research partnership led by Laura Stephenson at the University of Western Ontario. Data from C-Dem include two major Canadian Election Studies, four Provincial Election Studies, and four large-scale "Democracy Checkup" surveys. We also gathered data from the 2017 and 2018 Canadian Municipal Election Study (consisting of surveys in Calgary, London, Mississauga, Montreal, Quebec City, Toronto, Vancouver, and Winnipeg), the 2021 Calgary election study, and the 2022 Ontario Municipal Election Study. Finally, we assembled a large-scale survey in Ontario (led by Nathan Grasse) and a pan-Canadian survey (led by Jack Lucas, in partnership with the Samara Centre). In total, we amassed 94,073 survey responses available from 2,717 unique municipalities.

Table SM.1: Survey Data Sources for MRP Model

Name	Year	Geography	N	Project
Canadian Municipal Election Study	2018	Major Cities	12569	CMES
Canadian Election Study	2019	Canada	29582	C-Dem
Education Policy Attitudes Study	2019	Ontario	2913	Grasse et al.
Democracy Checkup	2020	Canada	1201	C-Dem
Provincial Election Study (BC)	2020	British Columbia	1466	C-Dem
Provincial Election Study (NB)	2020	New Brunswick	1042	C-Dem
Provincial Election Study (SK)	2020	Saskatchewan	975	C-Dem
Calgary Election Study	2021	Calgary	2050	CMES
Canadian Election Study	2021	Canada	17942	C-Dem
Provincial Election Study (NL)	2021	Newfoundland	810	C-Dem
Provincial Election Study (NS)	2021	Nova Scotia	1100	C-Dem
Samara Attitudes and Identities Survey	2021	Canada	3543	CMB
Democracy Checkup	2022	Canada	9583	C-Dem
Ontario Municipal Election Study	2022	Ontario	3897	MLDP
Provincial Election Study (ON)	2022	Ontario	3919	C-Dem
Provincial Election Study (QC)	2022	Quebec	1481	C-Dem

2 Survey Data (Politicians)

Table SM.2 describes the samples from the 2020–2024 Canadian Municipal Barometer (CMB) surveys with local politicians. This table includes each survey's overall response rate. It also compares each sample's observable characteristics (politicians' province, municipality size, and gender) with those of the population it came from.

Table SM.2: CMB Populations vs. Samples

	2020		**2021**		**2022**		**2023**		**2024**	
	Pop. %	Samp. %	Pop. %	Samp. %	Pop. %	Samp. %	Pop. %	Samp. %	Pop. %	Samp. %
Response Rate		22		21		23		26		21
Province										
AB	10	12.2	10	15	9	11	10	12	10	11
BC	12	12.4	12	14	11	9	12	13	12	14
MB	3	3.7	3	2	3	2	3	3	3	2
NB	3	2.5	3	2	3	3	3	2	3	3
NFLD	1.5	1.7	1.5	10	1.5	1	1	1	1	1
NS	3	3.4	3	6	3	3	3	3	3	5
NWT	0.2	0.1	0.2	1	0.2	0.1	0.1	0.01	0.1	0
ON	36	35.1	36	34	36	32	35	32	35	34
PEI	0.7	0.9	0.7	9	0.7	0.9	0.01	0.01	0.01	0.01
QC	28	25.4	28	22	29	35	29	21	29	23
SK	2.5	2.5	2.5	3	2	3	2	3	2	3
YT	0.2	0.3	0.2	2	0.2	0.2	0.2	0	0.2	0.1
Pop. Size										
9–15K	29	29.3	30	29	30	26	28	25	49	47
15–25K	21	19.6	21	19	21	19	22	18		
25–50K	16	16.5	16	16	16	15	17	15	17	17
50–100K	12	10	12	12	12	12	12	10	12	14
100–500K	14	16.2	14	16	14	15	14	13	14	15
500K+	8	8.5	7	8	7	8	7	8	7	8
Gender										
Men	68	69.5	68	66	65	41	64	60	64	62
Women	32	30.5	32	34	35	59	36	39	36	38

Note: In 2024, the population-size categories of 9–15K and 15–25K were merged.

3 MRP Model

As noted in the main text, we use Multilevel Regression and Poststratification (MRP) to construct our estimates of municipalities' average ideology scores. This requires two kinds of data: individual-level survey responses (discussed on page two of our Supplementary Material) and information on the sociodemographic composition of municipal residents. In this section, we provide additional information on the later data source, as well as validity tests for our MRP estimates.

3.1 Data on the Sociodemographic Composition of Municipal Residents

We collected data on the sociodemographic composition of municipal residents from two sources. First, we collected 79 unique sociodemographic indicators from the 2021 Canadian census, including information on income, racial composition, educational attainment, housing stock, population size, population density, age, immigration status, religion, language, occupation type, and commuting patterns. Second, we used data on Conservative Party support in Canada's 2021 federal election, using areal-weighted interpolation to interpolate election results at the polling-station level into municipalities.

3.2 MRP Model Specification

We model individual residents' left-right ideologies as follows:

$$y_i = \theta_0 + \alpha_{j[i]}^{age.sex.edu} + \alpha_{k[i]}^{mun} + \alpha_{l[i]}^{region}$$

In this model, y represents each individual i's self-placement on a 0–10 scale, θ represents an overall intercept, α_j captures varying individual-level demographic intercepts, α_k captures varying municipality-level intercepts, and α_l captures varying regional intercepts. We model the age, gender, and education intercepts as drawn from a normal distribution with mean zero:

$$\alpha_j^{age.sex.edu} \sim \mathcal{N}(0, \sigma^2_{age.sex.edu})$$

We model the municipal intercepts as predicted by the regional intercepts (that is, the region in which the municipality is situated) as well as a set of k municipality-level predictors, γ_k:

$$\begin{aligned} \alpha_k^{mun} &\sim \mathcal{N}(\mu_k^{mun}, \sigma_k^2) \\ \mu_k^{mun} &\sim \alpha_{l[i]}^{region} + \gamma_{k1} ... \gamma_{kn} \end{aligned} \tag{1}$$

We use Bayesian model averaging (BMA) to select well-performing predictors from the 79 census indicators as well as the municipality-level Conservative Party support variable described above, retaining all variables with a posterior model inclusion probability above 50%. This procedure indicated that the following variables should be included in the model: percentage in the bottom income decile, percentage English speakers, percentage White, percentage Christian, percentage with university degrees, and Conservative Party 2021 support.

Finally, we assume that the regional intercepts (BC, Prairies, Ontario, Quebec, Atlantic Canada) are drawn from a normal distribution with mean zero:

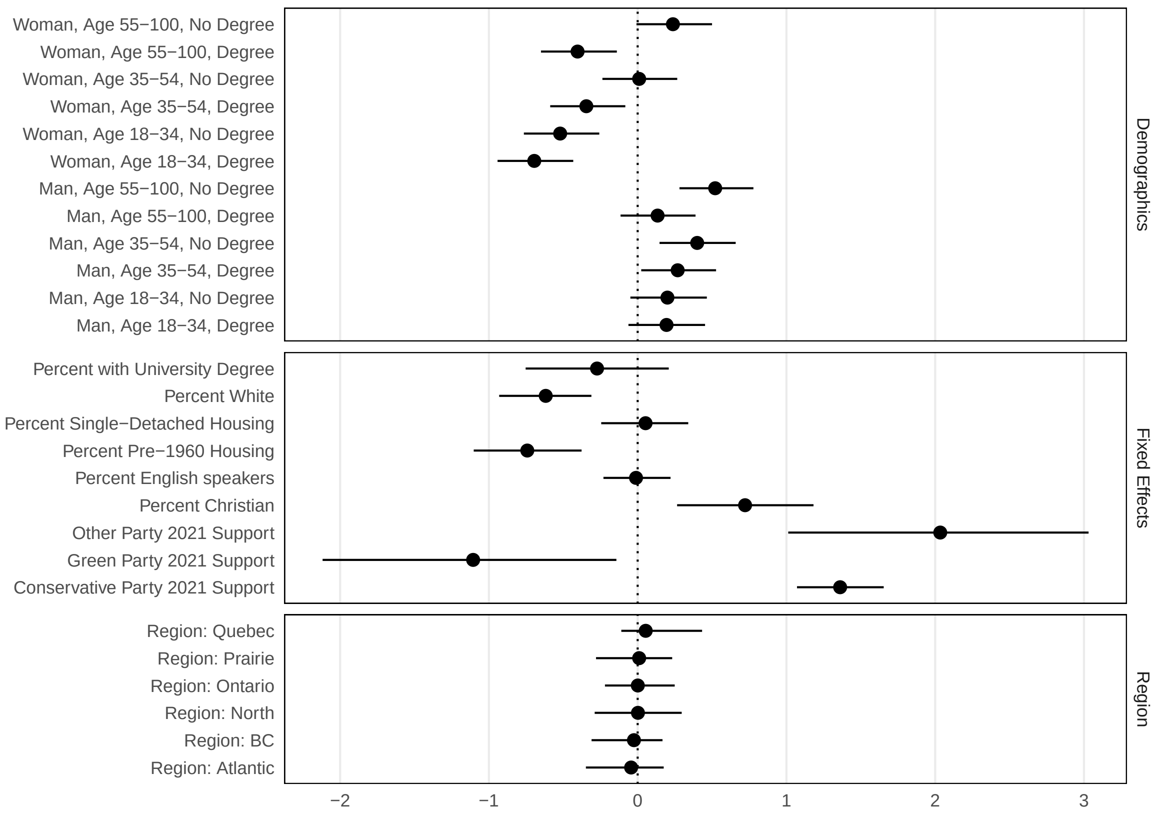

Figure SM.1: Fixed and Random effects from Multilevel Regression Model

$$\alpha_l^{region} \sim \mathcal{N}(0, \sigma^2_{region})$$

We use diffuse default priors for all γ parameters. We use Stan, as implemented in the rstanarm package in R (Goodrich et al., 2020), to fit our models, drawing 2,000 samples from each of four chains following a warm-up period of 2,000 iterations. Post-estimation tests provide strong evidence of convergence: $\hat{R}$ values are 1.0 for all parameters, and traceplots show clear evidence of mixing.

3.3 Model Results

To better understand the performance of the multilevel regression model, Figure SM.1 summarizes the fixed and random effects from the model. (Note that the 1,769 unique municipal intercepts were also included in the model but are not visualized here.) In general, the demographic intercepts conform to our expectations: Women are generally to the left of men (in keeping with ideological gender gaps), those with university degrees are to the left of those without university degrees (in keeping with ideological education gaps), and younger respondents are to the left of older respondents, especially among women (in keeping with ideological age gaps). Municipality-level predictors also align with our expectations and are generally high-precision coefficients. Finally, conditional on the other variables in the model, regional intercepts vary more modestly. In general, the results in Figure SM.1 suggest that the model performs as expected—and well—in predicting residents' ideological self-placements.

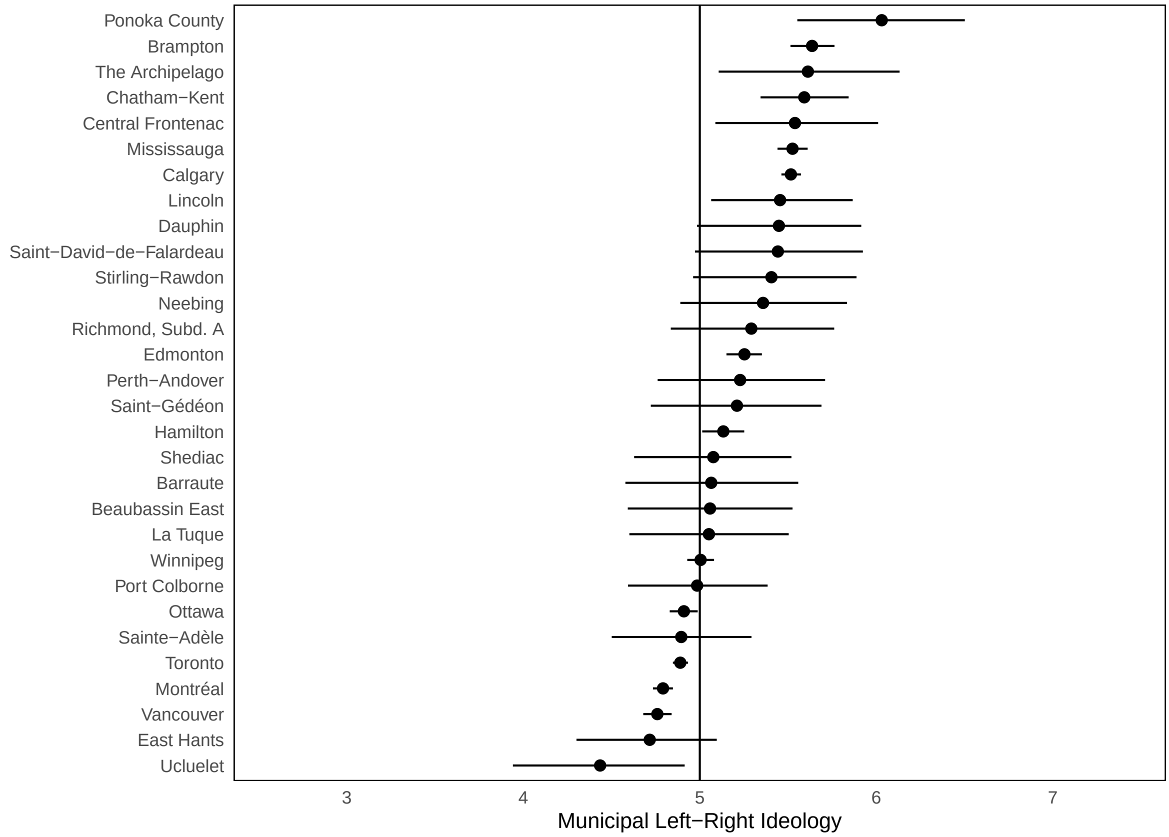

Figure SM.2: Municipal MRP Estimates: Sample of Results

3.4 MRP Estimates and Uncertainty

To develop targets for post-stratification, we used data from the 2021 Canadian census to calculate the proportion of each municipality in each age-sex-education combination. We then used these demographic categories, along with the municipal predictors, to estimate average ideological self-placement scores for each municipality. We provide a sample of our results in Figure SM.2, which contains the minimum and maximum estimates, the ten largest municipalities in Canada, and a random sample of twenty additional municipalities. Points in the figure represent MRP point estimates, and whiskers are 95% credible intervals.

The results in Figure SM.2 have good face validity. Municipalities known to be more left-leaning—such as Ucluelet (a municipality on Vancouver Island), East Hants, and Vancouver —are at the bottom of the figure. Likewise, the most conservative municipality—Ponoka County, in rural Alberta—comports with expectations. Notice that the precision of the municipal estimates clearly varies when we compare the large municipalities for which more survey responses are available (e.g., Calgary) to smaller municipalities (e.g., Ponoka County). Even so, our estimates, even for very small municipalities, are quite precise, allowing for meaningful differentiation between municipalities.

Clearly, however, the credible intervals in Figure SM.2 illustrate an often-ignored feature of the MRP estimates: While our very large database of survey responses allows for unprecedentedly precise estimates, there is still meaningful *uncertainty* about the average ideological self-placement of each municipality. While many MRP-based studies ignore this uncertainty, findings in both the broader measurement literature (Treier and

Jackman, 2008) and in more recent MRP-based research (Lucas, 2024; Tausanovitch and Warshaw, 2014) have illustrated the value of a more conservative approach that incorporates measurement uncertainty into subsequent estimates. To do so, we adapt the standard Monte Carlo Integration procedure to suit MRP, as follows:

1. Using posterior draws from the fully Bayesian MRP model, calculate 1,000 plausible MRP estimates of average ideological self-placement for each municipality.

2. Using an intercept-only model ($y = \alpha$), estimate the average perceptual error among politicians for each of the 1,000 plausible MRP estimates.

3. Draw one value from the normal posterior distribution of each α [$\alpha \sim \mathcal{N}(0, \sigma^2)$], where α is the intercept value and σ^2 is the standard error of the intercept.

4. Summarize point estimates and credible intervals using the 2.5^{th}, 50^{th}, and 97.5^{th} percentiles of this vector of 1,000 values.

4 Perceptual Errors: Constituent Ideology vs. Policy Preferences

Our analysis in the main text focuses on politicians' perceptions of their constituents' positions on the left-right ideological spectrum. However, many analyses of politicians' perceptual accuracy and conservative bias have focused on citizens' policy preferences. Here, we demonstrate that perceptual errors about constituent ideology are closely related to perceptual errors about constituents' policy preferences.

To do this, we utilize data from the 2020 Canadian Municipal Barometer survey, which asked politicians to estimate their constituents' attitudes about various policy issues. We use survey data from the Digital Democracy Project to estimate the *actual* average attitude of residents in each municipality, applying multilevel regression and poststratification (MRP) to measure public opinion.

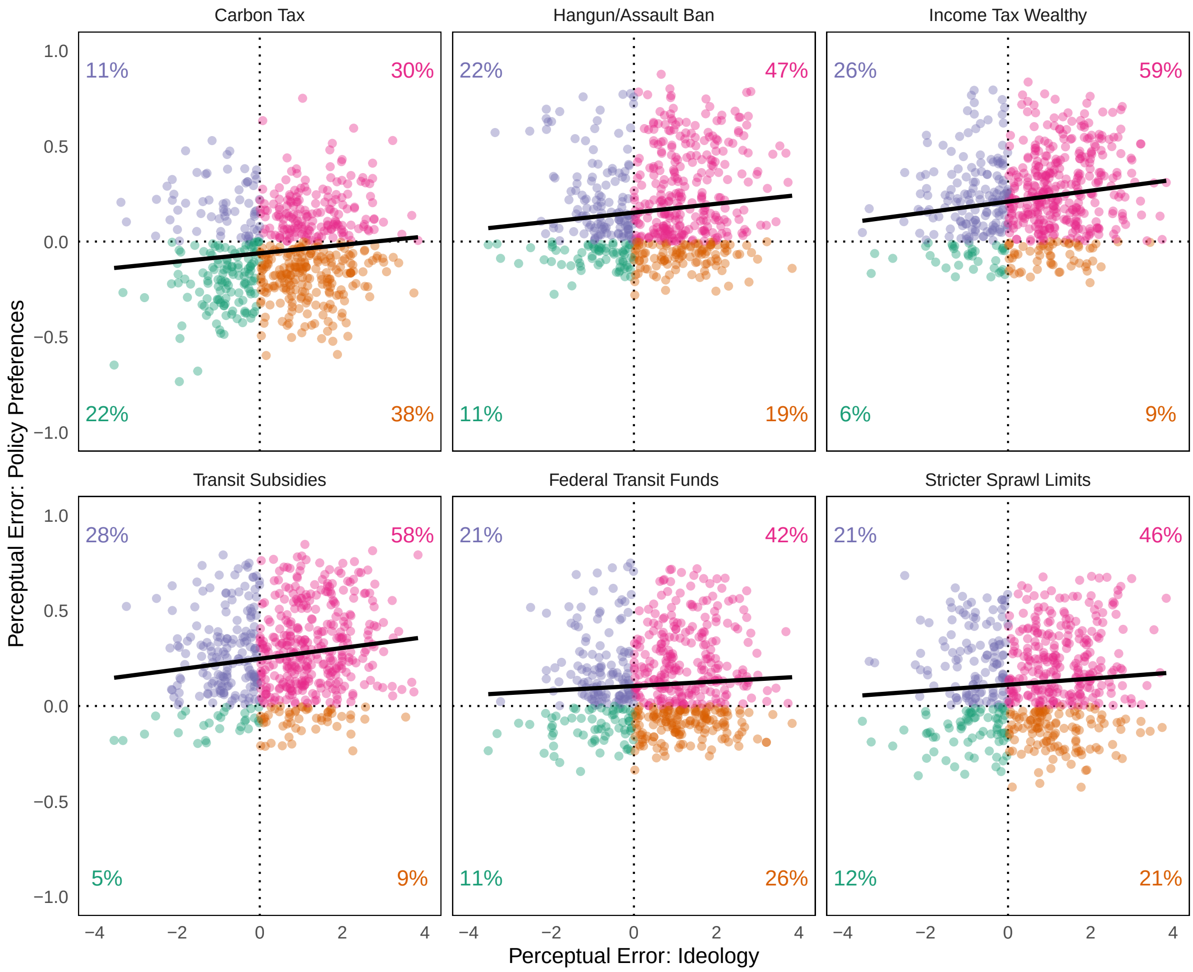

Figure SM.3: Perceptual Errors about Constituent Ideology versus Policy Preferences

Figure SM.3 provides a summary of our validity test. Each panel in the figure focuses on a specific policy issue. Along the horizontal axis, we report politicians' perceptual errors about their average constituent's ideology—the quantity that is our focus in the main text. Along the vertical axis, we report politicians' perceptual errors about their average constituent's policy preference. Positive values indicate that a politician has over-estimated their constituents' conservatism, whereas negative values indicate politicians have under-estimated their constituents' conservatism.

As shown in Figure SM.3, for all policy issues, perceptual errors about constituent ideology are positively related to perceptual errors about constituents' policy preferences. Moreover, for five of the six issues, the plurality of politicians over-estimate their constituents' ideological and policy conservatism. In short, our descriptive results in the main text do not depend on our decision to focus on ideology rather than policy preferences.

5 Question Wording and Order

Figures SM.4, SM.5, and SM.6 provide screenshots of the questions answered by politicians to estimate the distribution of their constituents' ideologies (perceived-distribution task), to estimate the ideology of their average constituent (point-estimate question), and to indicate their own ideology (ideological self-placement). Note that politicians were randomly assigned to one of four possible question orderings:

1. Perceived-distribution task, point-estimate question, ideological self-placement
2. Point-estimate question, perceived-distribution task, ideological self-placement
3. Ideological self-placement, perceived-distribution task, point-estimate question
4. Ideological self-placement, point-estimate question, perceived-distribution task

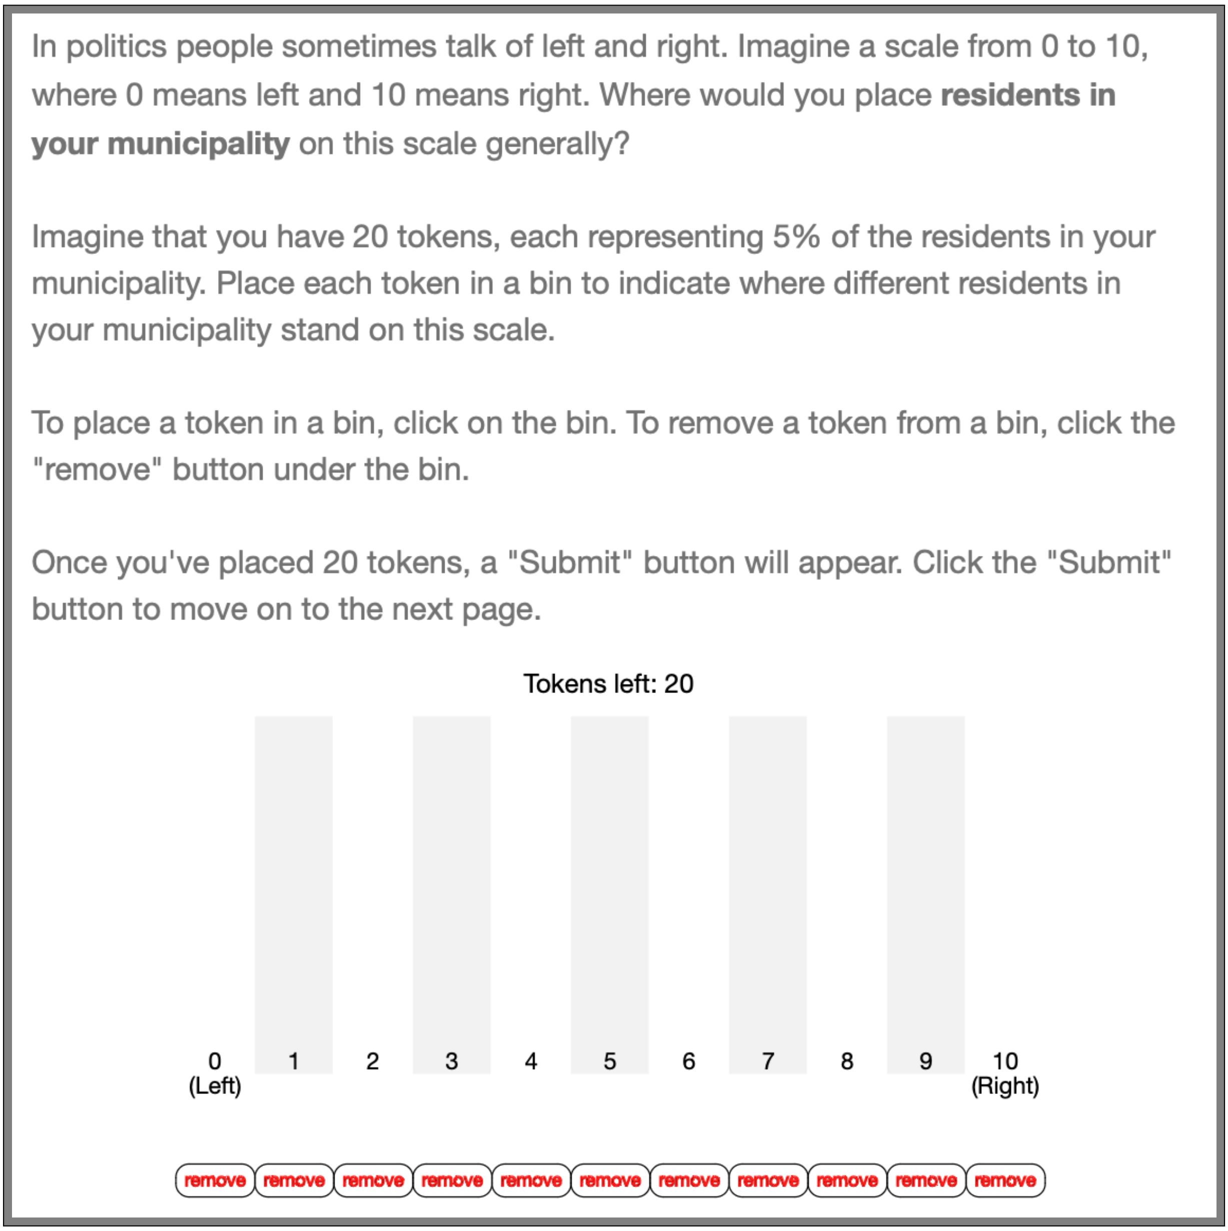

Figure SM.4: Perceived-Distribution Task

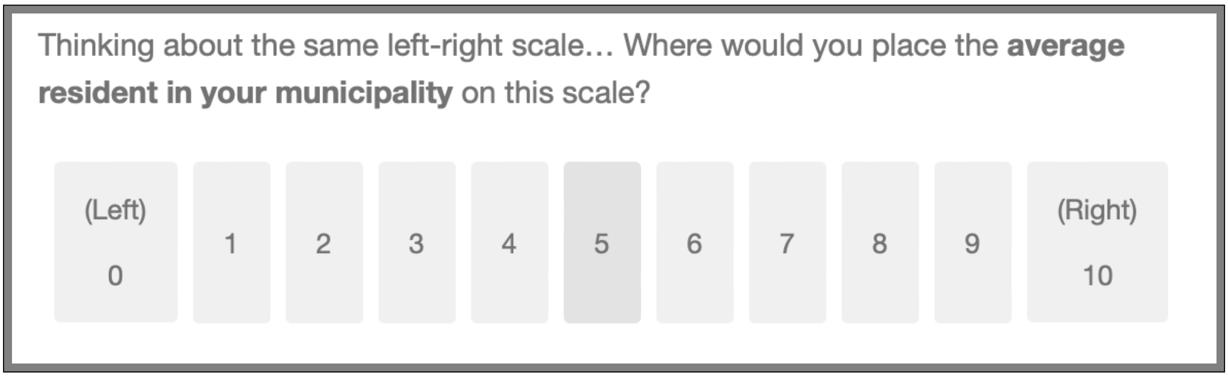

Figure SM.5: Point-Estimate Question

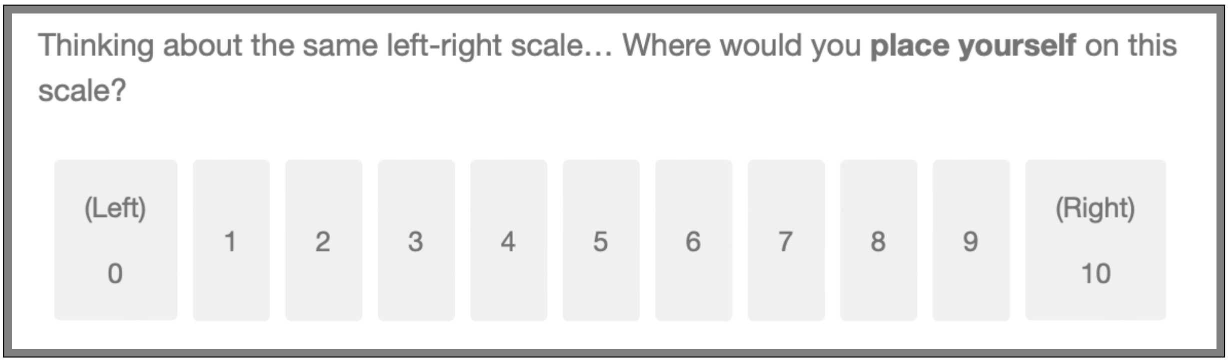

Figure SM.6: Self-Reported Ideology

6 Implementing the Perceived-Distribution Task

To make the perceived-distribution task easy to implement in Qualtrics, we have included a `.qsf` file in our replication materials. One can also implement the perceived-distribution task in Qualtrics by following the steps below. See hyperlinks for illustrative images.

One, create a *Text/Graphic* question. Two, within that question, access the *Rich Content Editor*. Three, click the *Source* button. Four, copy-and-paste the HTML code below. This code creates a space on the question page for the perceived-distribution task:

```
<div style="display: inline-block; position: relative; width:100%;
padding-bottom: 50%; vertical-align: top; overflow:hidden"
id="dist_body"> </div>
```

(Note: Text can be added above this space-holder.) Five, having selected the same question, look to the left of the screen and click the *JavaScript* button within the *Builder* menu. Six, copy the contents of the `.js` file in our replication materials, and paste these contents into the *addOnReady* function within the JavaScript window you opened in Qualtrics. This code generates the perceived-distribution task and records participants' responses. Seven, click the *Save* button at the bottom-right of the window. Eight, look to the left of the screen and click on *Survey Flow* button. Nine, add an *Embedded Data* element to the top of your *Survey Flow.* Ten, add `dist_response` and `dist_response_order` as variables within the *Embedded Data* element. Our JavaScript code will record participants' responses to the perceived-distribution task within these variables. Eleven, click the *Apply* button at the bottom-right of the window. Twelve, look to the left of the screen and click on *Look and Feel* button. Thirteen, within that menu, select *General* sub-menu. Fourteen, within the *Header* field, paste the HTML code below. This code loads a JavaScript library called D3, which we need to generate the perceived-distribution task:

```
<script src="https://d3js.org/d3.v4.js"></script>
```

Thirteen, click the *Apply* button at the bottom-right of the screen.

You are done! When you download your data from your Qualtrics survey, participants' responses to the perceived-distribution task should be recorded within columns titled `dist_response` and `dist_response_order`. `dist_response` indicates how many tokens are in each of the eleven bins (e.g., "3,3,3,4,4,3,0,0,0,0,0"). `dist_response_order` provides a participant's exact sequence of token additions and deletions (e.g., "ADDbin-1,DELbin-1,ADDbin-5...").

7 Additional Validation of the Perceived-Distribution Task

The perceived-distribution has previously been validated using data from American citizens (Dias, Lelkes and Pearl, 2024). That said, we perform additional analyses to ensure that politicians' responses to the perceived-distribution task are not mechanistic. The results of these analyses are displayed in Figure SM.7.

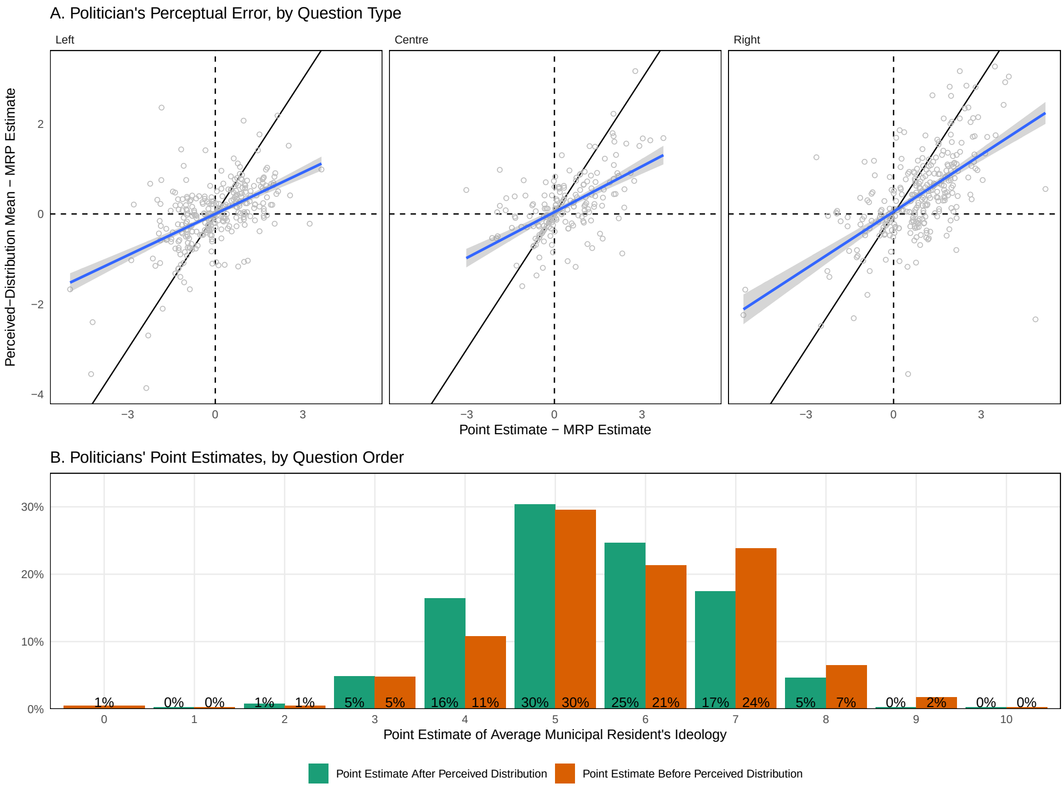

Figure SM.7: Panel A: Relationship between perceptual error on point-estimate question (horizontal axis) and perceptual error on the perceived-distribution task (vertical axis). Panel B: Distribution of point estimates by question order.

Panel A plots perceptual-error scores from the point-estimate question (horizontal axis) against perceptual-error scores from the perceived-distribution task (vertical axis). We summarize the relationship between the two sets of perceptual-error scores with the blue line, which is drawn from an outlier-robust bivariate regression model. If the error on the two tasks were perfectly correlated, the blue line would follow the dark black 45-degree line, indicating that the perceived-distribution task offered no improvement over the point-estimate question. In fact, however, the blue line is much flatter than the 45-degree line. That said, perceptual error on the point-estimate question is still positively correlated with perceptual error on the perceived-distribution task. This finding confirms that the perceived-distribution task reduces perceptual error in proportion to the size of these errors.

In Panel B, we report the distribution of point estimates among politicians who completed the point-estimate question before the perceived-distribution task (orange bars) and those who completed the point-estimate question after the task (green bars). These distributions indicate that the most common effect of the perceived-distribution task is to shift politicians' point estimates from the moderate-to-strong right of the spectrum (positions 7 and 8) toward the moderate right (6) and moderate left (4). Interestingly, the effect of the question-order experiment on politicians' point estimates is neither to substantially alter choices at the extremes nor to dramatically increase the proportion of politicians who select the scale midpoint (5).

References

Dias, Nicholas C, Yphtach Lelkes and Jacob Pearl. 2024. "American partisans vastly under-estimate the diversity of other partisans' policy attitudes." *Polit. Sci. Res. Meth.* pp. 1–11.

Goodrich, Ben, Jonah Gabry, Imad Ali and Sam Brilleman. 2020. "rstanarm: Bayesian applied regression modeling via Stan.". R package version 2.21.1.
URL: *https://mc-stan.org/rstanarm*

Lucas, Jack. 2024. *Ideology in Canadian Municipal Politics.* Toronto: University of Toronto Press.

Tausanovitch, Chris and Christopher Warshaw. 2014. "Representation in Municipal Government." *American Political Science Review* 108(03):605–641.

Treier, Shawn and Simon Jackman. 2008. "Democracy as a Latent Variable." *American Journal of Political Science* 52(1):201–217. https://onlinelibrary.wiley.com/doi/10.1111/j.1540-5907.2007.00308.x.