Online Appendix:
How Should We Measure District-Level
Public Opinion on Individual Issues?

1. Introduction

In this appendix, we address three issues that space considerations precluded from including in the main body of our article. First, we show the relationship between the predictors in our model and public opinion on same-sex marriage. This helps demonstrate the systematic relationships we find between demographic and geographic predictors and public opinion. Second, we show how the gains from MRP are related to the sample size in a particular district. Third, we examine how well MRP predicts the public opinion of voters rather than all adults.

2. What is the relationship of particular demographic and geographic predictors with public opinion?

Our main article does not describe the specific relationship between public opinion and particular demographic and geographic predictors in our model. However, the strong performance of MRP estimates is built on strong and predictable relationships between individual- and district-level demographic predictors and public opinion.

In Figure A-1, we show the coefficients from our full multilevel logistic regression and poststratification (MRP) model of public opinion on a federal same-sex marriage amendment. The survey responses are coded 1 for support of an amendment to ban same-sex marriage and 0 otherwise ("no," "don't know," or "refused"). Thus, positive coefficients in the figure indicate greater support of an amendment to ban same-sex marriage. The results shown in the figure are based on a national sample of 5,000 respondents from the 2004 NAES, 2006 CCES, and 2008 CCES.

The figure shows that standard demographic predictors perform reasonably well. First, attitudes on same-sex marriage are strongly linked with gender. Second, higher levels of education are correlated with more liberal attitudes on same-sex marriage. Perhaps surprisingly, race is not correlated with same-sex marriage attitudes after controlling for other factors. However, several district-level predictors are strongly related to same-sex marriage attitudes. Respondents in wealthier districts generally have more liberal attitudes on same-sex marriage. Respondents in districts with higher proportions of same-sex couples also have more liberal attitudes. Finally, respondents in states with higher percentages of Evangelicals and Mormons have more conservative attitudes.

3. What is the relationship between the sample size in a particular district and the accuracy of MRP estimates?

Our main article shows how MRP outperforms disaggregation at national samples between 2,500 and 30,000. But it does not show whether MRP outperforms disaggregation across congressional districts with a variety of sample sizes. In Figure A-2, we show that MRP dramatically outperforms disaggregation in districts with smaller samples. If there are less than 20 respondents in a district, MRP yields roughly 50% better estimates than disaggregation. The gap closes in districts with larger samples, and the performance of MRP and disaggregation essentially converges in districts with more than 100 respondents. Our findings are similar to Lax and Phillips' (2009) state-level analysis, which finds that MRP outperforms disaggregation in small- and medium-sized states, but their performance converges in high-population states (Lax and Phillips 2009, Figure 1). Future research should further analyze whether MRP continues to outperform

disaggregation in larger district-level samples.

4. How well does MRP predict the public opinion of voters?

In the main body of our article, we briefly address the concern that because of the unavailability of crucial district-level data broken down by voters and nonvoters, our MRP estimates should be understood as estimates of the opinions of adults rather than those of the smaller pool of actual voters. Yet Park, Gelman, and Bafumi (2004) find that MRP can produce excellent estimates of state presidential election outcomes,[1] and our article presents relatively high correlations between MRP estimates and referendum results.

Moving beyond these validations, here we seek additional reassurance that researchers might profitably use MRP estimates of adult public opinion as proxies for the opinions of voters. We perform additional split-sample analyses comparing our estimates to "baseline" estimates of voters using CCES respondents that self-reported that they voted.[2] In Figure A-3, we show that that MRP estimates yield very similar correlations with the baseline public opinion of voters and all adults.[3] On average, the correlations of the MRP estimates with the baseline public opinion of voters are just .00-.09 lower than

---

[1] Park, Gelman, Bafumi (2004) estimate an MRP model of state-level vote intention in the 1988 and 1992 presidential races using preelection polls poststratified by state, region, and several demographic variables. Their model includes just 2200 survey respondents in 1988 and 4650 survey respondents in 1992. Despite these relatively small sample sizes, they find that the MRP estimates are extremely good predictors of the presidential election results. The mean absolute errors of the state estimates produced by the MRP model is just 4 percentage points, compared to between 5 and 10 percentage points for simpler models.

[2] For this analysis, we generated 20 simulations of 5,000 people to evaluate the average correlation of our MRP estimates with baseline estimate of public opinion using only CCES respondents that indicated they voted.

[3] Some of the small differences we found could be due to the fact that only the CCES includes a postelection survey. As a result, we are not able to include NAES respondents in our baseline estimate of the public opinion of voters.

the correlations with all adults. This is encouraging, but further efforts to distinguish

between voters and nonvoters would be a worthwhile endeavor for future refinements of

MRP techniques. Above all, such efforts will require researchers to link individual-level

data on turnout with census microdata.

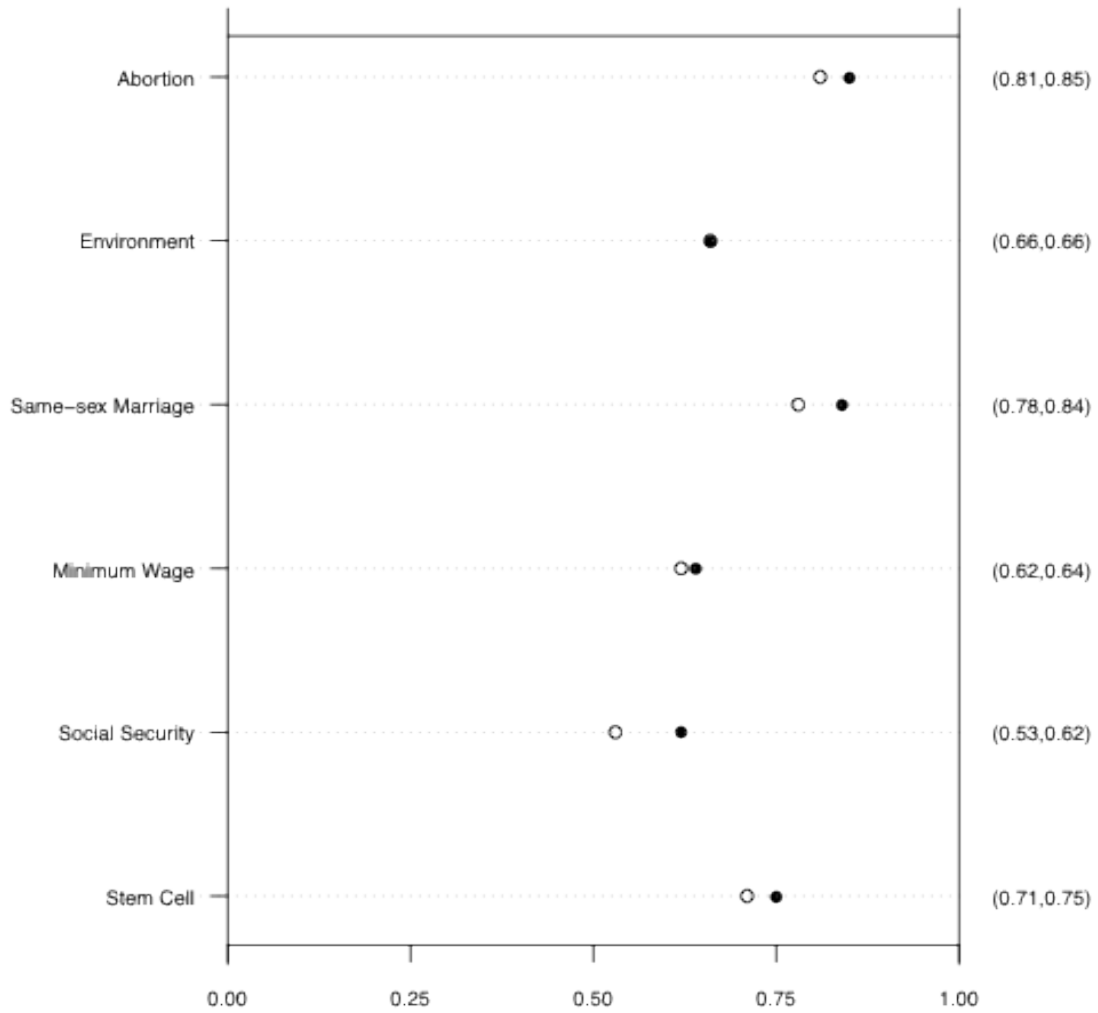**FIGURE A–1: Estimates of Hyperparameters for Gay Marriage Amendment**
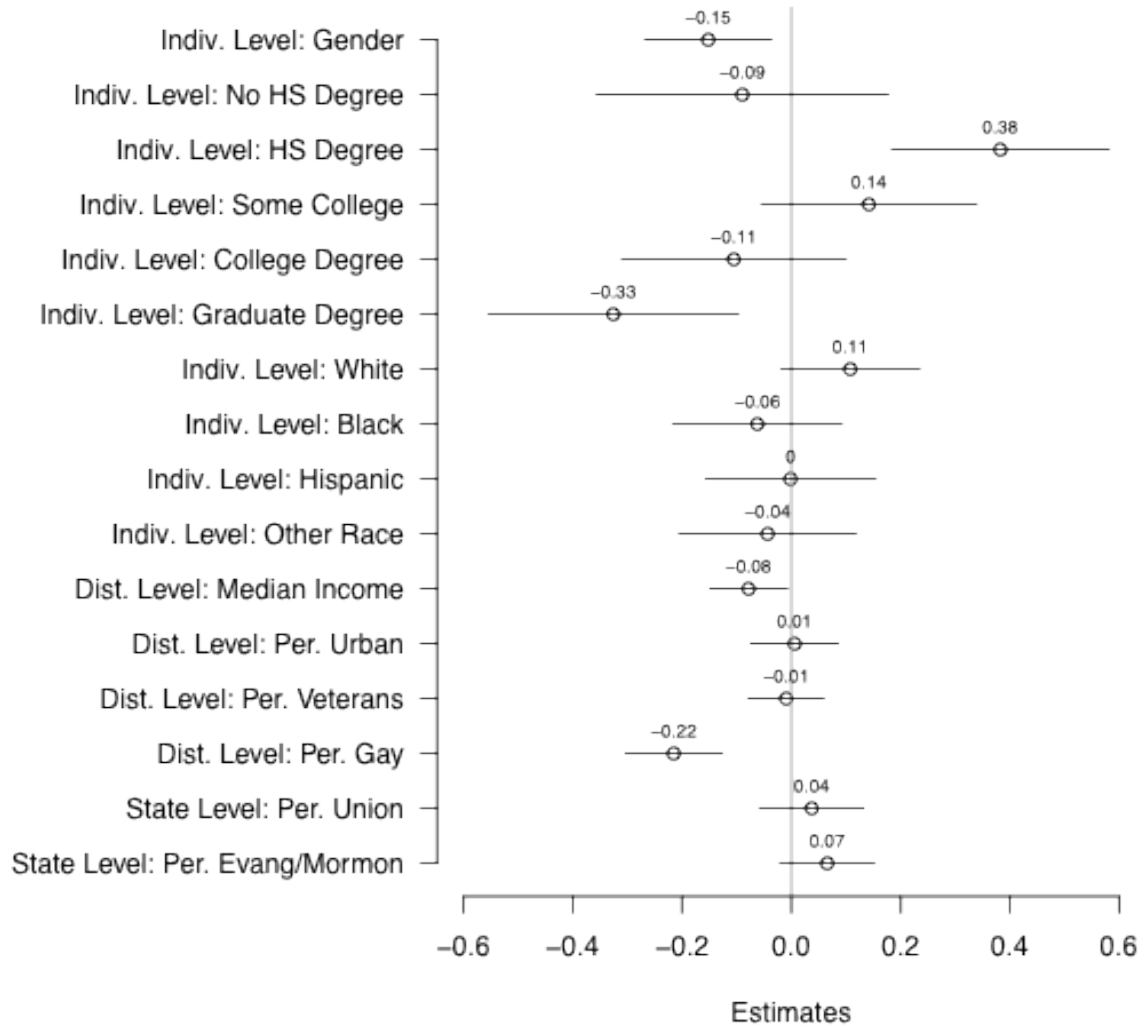
This figure shows estimates and 95% credible intervals for the geographic and demographic coefficients in our MRP logistic regression of the probability of supporting a federal amendment to ban same-sex marriage.

FIGURE A–2 Cross Validation –– Mean Errors for Same–Sex Marriage by District Sample Size
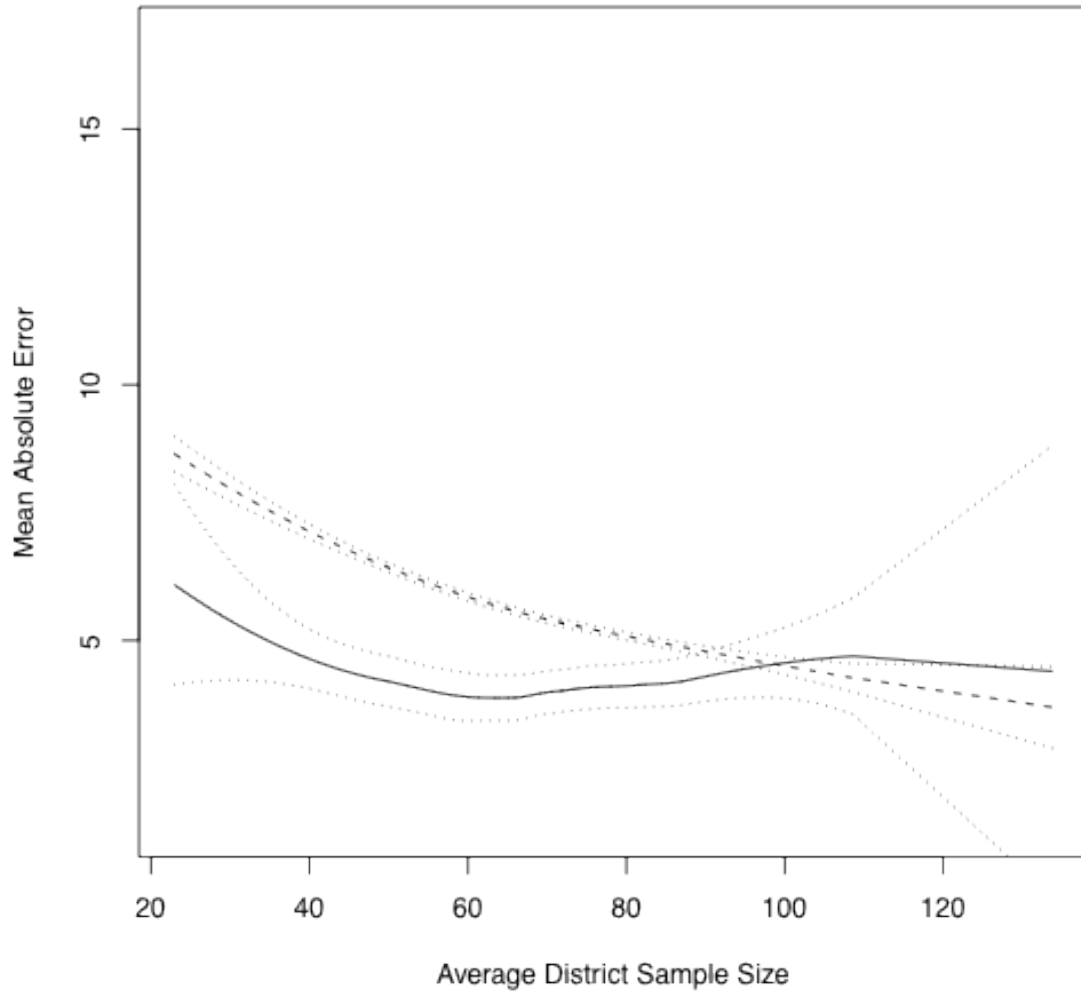
This figure uses lowess curves, surrounded by 95% confidence intervals, to plot the mean absolute error by district against the sample size in that district using MRP (the solid curve) and disaggregation (the dashed curve). The results are based on 50 simulations of national samples with 30,000 respondents.

FIGURE A–3: Comparing Model Accuracy for Voters v Non–Voters

Abortion	(0.81,0.85)
Environment	(0.66,0.66)
Same–sex Marriage	(0.78,0.84)
Minimum Wage	(0.62,0.64)
Social Security	(0.53,0.62)
Stem Cell	(0.71,0.75)

0.00   0.25   0.50   0.75   1.00

This figure compares the correlation of MRP estimates with voters and all adults for six issues based on 20 simulations with national samples of 5,000 respondents. For each issue, it plots the correlations of MRP estimates with baseline estimates of the public opinion of self-reported voters (∘) and all adults (•) in each congressional district.