Subnational Public Opinion Estimation Using MrsP

Joseph T. Ornstein*

September 10, 2017

Abstract

Multilevel regression and poststratification (MrP) is a powerful tool for estimating local public opinion from national survey data. However, it has seen little use outside of the United States, owing to its rather stringent data requirements. In this paper, I prove that an extension of MrP using synthetic poststratification (MrsP) can overcome these limitations without introducing error. I then use this technique to disaggregate cross-national survey data from the World Values Survey and International Social Survey Programme. I find that these estimates are highly correlated with alternative measures of public opinion, and discuss how this technique could serve as a useful tool for election forensics.

1 Introduction

There are many research questions in comparative politics that would benefit from detailed local public opinion data. For example, scholars of Chinese politics might be interested to know where pro-democracy sentiment is the strongest. Is Hong Kong an outlier, or do the residents of Shanghai and Chongqing share similar values? Scholars studying civil conflict may want to identify local areas where citizens are most likely to support secession. And researchers who conduct election forensics might like to compare local public opinion data against observed vote outcomes. (Did 99.5% of Chechnyans really support Vladimir Putin in the 2012 Russian presidential election?)

Unfortunately, we rarely have access to such detailed local public opinion data. In many countries, there are few national public opinion surveys, and fewer still that include enough observations to make meaningful local level inferences. One approach to this problem, pioneered by Park et al. (2004), is multilevel regression and poststratification (MrP). This technique produces subnational public opinion estimates by first fitting a multilevel regression model, then reweighting the model's predictions according to each local unit's observed demographic characteristics – a process called poststratification. By producing local area estimates where we previously lacked detailed polling information, MrP has opened up a

*PhD Candidate, University of Michigan. Many thanks to Paul Kellstedt, Rob Franzese, and the participants at MPSA, the Michigan Political Economy Workshop, and the Statistical Learning Workshop for their helpful comments and suggestions.

wide range of research questions for study. Recent applications include the study of US state-level support for same-sex marriage (Lax & Phillips 2009), representation in US state politics (Lax & Phillips 2012), and US municipal politics (Tausanovitch & Warshaw 2014).

However, MrP has yet to see extensive use outside of the United States, owing to its rather stringent data requirements. Because the poststratification stage requires a joint distribution across every demographic variable included in the regression model, it is only possible to generate estimates in countries that provide detailed public census crosstabulations or local level microdata. If we are interested in estimating subnational opinion for countries without such detailed census data, what is the proper way forward?

Leemann & Wasserfallen (2016) introduce a promising refinement to MrP, which they call multilevel regression and *synthetic* poststratification (MrsP). Rather than creating poststratified estimates using the true joint distribution of the demographic variables in the individual-level model, this approach proceeds as if the demographic variables are statistically independent. Then, the poststratification weights can be derived from the product of the marginal distributions, a process they call synthetic poststratification. The authors conduct a Monte Carlo test of this procedure, demonstrating that, so long as the demographic variables are not too strongly correlated with one another, MrsP estimates do not significantly diverge from those of classical MrP.

In this paper, I argue that Leeman & Wasserfallen, if anything, understate the case for adopting MrsP in comparative politics research. In the next section, I present a proof that MrP and MrsP produce identical estimates if the underlying individual-level model is linear-additive. This suggests that MrsP will produce *strictly* better estimates than MrP, by enabling the researcher to include additional individual-level predictors. In subsequent sections, I provide three empirical applications that illustrate the strengths of this new approach: (1) a study of UK public opinion on the European Union prior to the 2016 Brexit referendum, (2) a disaggregation of the World Values Survey's Secular-Rational versus Traditional values at the US state level, and (3) an examination of South African support for the African National Congress party at the election ward level, compared with results from the 2014 provincial elections. In each case, I show that estimates from MrsP correlate strongly with other subnational measures of public opinion, such as district-level vote shares. I conclude with suggestions for further research.

2 Proof: MrsP with Linear-Additive Models

Classical MrP produces local public opinion estimates through a two-step procedure. First, estimate an individual-level regression model from national-level survey data, where the modeled public opinion (y), is a function of some set of demographic and geographic characteristics. This model yields predictions for each category of respondent: $\hat{y} = f(X_{dem}, X_{geo})$. Second, we poststratify to each local area by taking the average prediction, weighted by the proportion of an area's residents in each category. If $\hat{\mathbf{y}}$ is the vector of predictions for each type of respondent, and $\mathbf{p}$ is the vector of observed probabilities of each type, then the poststratified estimate is the dot-product $\hat{\mathbf{y}} \cdot \mathbf{p}$. MrsP uses the same vector of predictions $\hat{\mathbf{y}}$, but uses a synthetic joint probability distribution, where each entry is the product of marginal probabilities. I will denote this synthetic poststratification vector as $\boldsymbol{\pi}$. Therefore, the poststratified MrsP estimates will be $\hat{\mathbf{y}} \cdot \boldsymbol{\pi}$.

Let X_1 through X_m be discrete random variables, and the $c \times m$ matrix X be a matrix in which the each row is one of the c possible combinations of values that X_1 through X_m can take. Crucially, we are not assuming that X_1 through X_m are independent, so $P(X_1 = x_{1i}, ..., X_m = x_{mk})$ need not equal $P(X_1 = x_{1i})...P(X_m = x_{mk})$.

Suppose the model is linear additive, such that $\hat{y} = X\hat{\beta}$. The vector of MrsP predictions for each unit is therefore $\pi'\hat{y}$, where π is the synthetic distribution vector. To complete the proof, we need to show that $\pi' X\hat{\beta} = p' X\hat{\beta}$. Because β is a vector, this is equivalent to showing that $\pi' X = p' X$.

$$p'X = \begin{bmatrix} \sum_i ... \sum_k P(X_1 = x_{1i}, ..., X_m = x_{mk}) x_{1i} \\ \vdots \\ \sum_i ... \sum_k P(X_1 = x_{1i}, ..., X_m = x_{mk}) x_{mk} \end{bmatrix}$$

$$= \begin{bmatrix} \sum_i P(X_1 = x_{1i}) x_{1i} \\ \vdots \\ \sum_k P(X_m = x_{mk}) x_{mk} \end{bmatrix}$$

$$= \begin{bmatrix} \sum_i ... \sum_k P(X_1 = x_{1i})...P(X_m = x_{mk}) x_{1i} \\ \vdots \\ \sum_i ... \sum_k P(X_1 = x_{1i})...P(X_m = x_{mk}) x_{mk} \end{bmatrix} = \pi' X$$

This completes the proof. If our underlying first-stage model is linear-additive, then our poststratified estimates will be identical whether we use MrsP or classical MrP. This frees us to include many more useful predictor variables in the first-stage regression, improving the precision of our estimates. In the following sections, I apply this technique to three empirical applications in comparative politics.

3 Brexit Referendum 2016

Following the surprising results of the 2016 UK Brexit referendum, many observers noted sharp regional differences in support for the measure. Areas with a more elderly population, and a smaller share of university-educated residents, were more likely to vote in favor of Brexit. Could these regional differences have been anticipated in advance?

To address this question, I generate MrsP estimates of local support for the European Union, taking as my survey data the 2015 International Social Survey Programme (ISSP) National Identity Module. The 2015 ISSP module asked a series of questions regarding opinion

Table 1: British Public Opinion Towards the European Union (ISSP 2015)

Question	**Agree**	**Neutral**	**Disagree**
"The UK should follow decisions of the European Union"	108	288	508
"The UK benefits from being a member of the European Union"	399	182	323
"I feel close to Europe"	195	70	639
"I would vote in favor of a referendum to remain an EU member"	300	285	319

toward the European Union. 35% of British respondents said they would vote against a referendum to remain the EU, compared to 33% who said they would vote to stay, and 32% who either could not choose, did not know what the EU was, or refused to answer. From responses to five questions regarding the European Union (see Table 3), I construct a measure of favorable opinion towards the EU using principal component analysis (taking the first principal component as my measure). By generating this continuous measure, I utilize more information than is contained in the binary "would you vote in favor of a referendum to remain in the EU?" survey question. I also create a more appropriate dependent variable for OLS regression, and can therefore leverage the results of my proof in section 2.

After generating the measure of EU support, I regress it on three demographic variables: age, education level, and marital status. The results from this regression are reported in Table 2.

Table 2: Coefficient estimates, regressing the EU Opinion variable on three demographic variables.

	Dependent variable:
	EU Opinion
Age 30 - 44	-0.097^{*}
	(0.053)
Age 45 - 64	-0.104^{**}
	(0.051)
Age 65 -	-0.105^{*}
	(0.055)
Secondary Qualifications	0.132^{***}
	(0.044)
University or Higher	0.418^{***}
	(0.052)
Married	-0.075^{**}
	(0.032)
Constant	-0.047
	(0.058)
Observations	904
R^2	0.101
Note:	*p<0.1; **p<0.05; ***p<0.01

With the predicted values from this regression in hand, I then generate poststratified es-

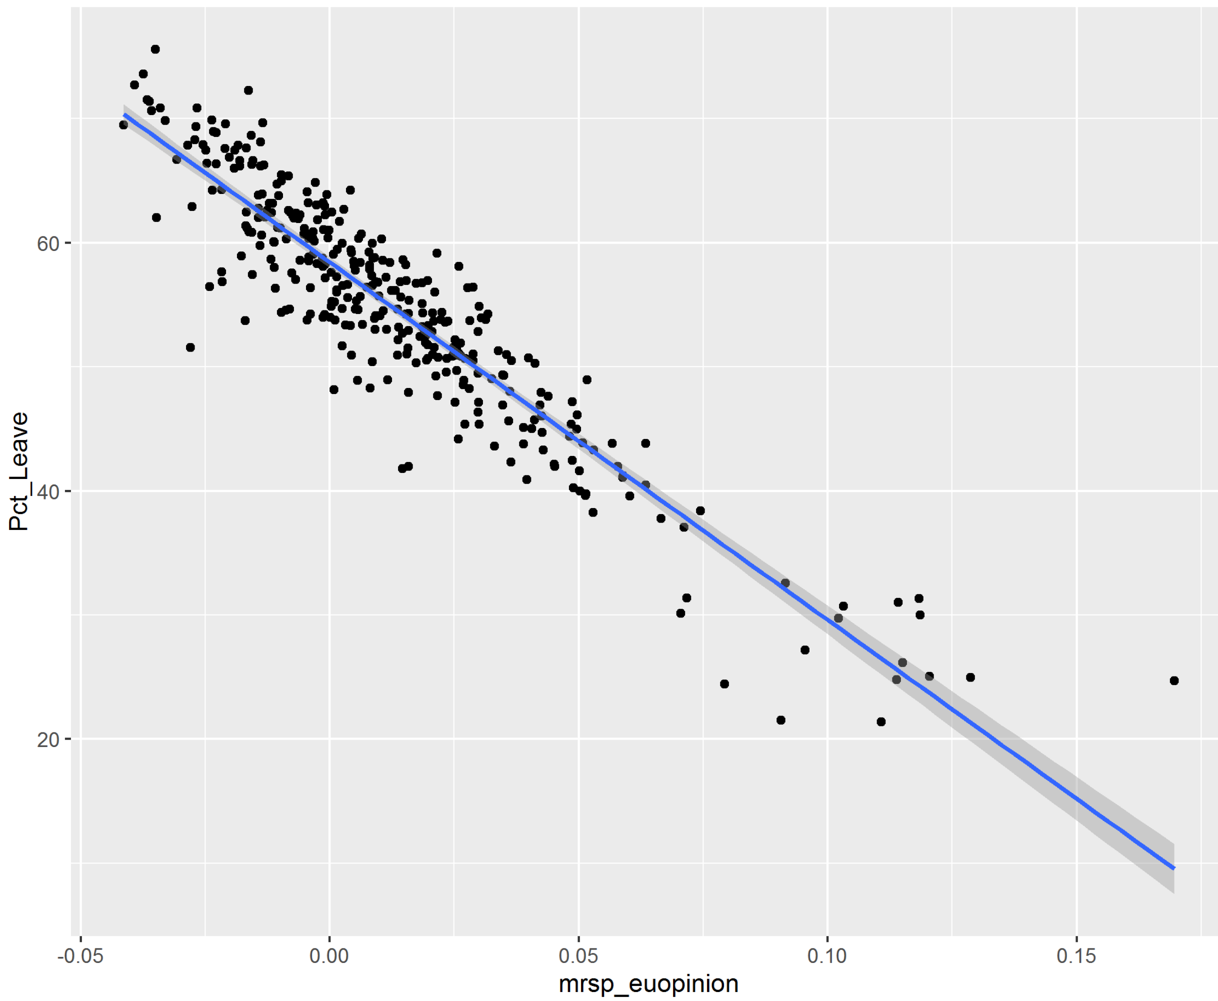

Figure 1: MrsP EU Opinion and District-Level Support for Brexit Referendum

timates of average local EU opinion, using a synthetic poststratification weights matrix from the 2011 UK Census (at the local authority level, the lowest level for which I have polling data from the Brexit referendum). As illustrated in Figure 1, the correlation between the MrsP estimates and Brexit results is striking. The City of London, however, is a notable outlier (plotted on the far right hand side of the figure). That area saw about fifteen percentage points more support for Brexit than the MrsP estimate would suggest. (This is particularly odd, given the concentration of the financial industry in the City of London. I would have thought, if anything, that MrsP would overestimate that area's support for Brexit.)

4 World Values Survey

The World Values Survey (WVS) is an invaluable source of comparative public opinion data (WVS, 2016). One product of this survey research is the so-called "Cultural Map of the World", which compares nations by their respondents' expressed Traditional vs. Secular-

Rational and Survival vs. Self-Expression values. This map yields interesting and consistent patterns over time, but it obscures substantial within-country variation on these measures. For instance, the US South is likely to express far more traditional values than other regions of the country. Could MrsP shed light on this regional heterogeneity?

To investigate, I construct US state-level MrsP estimates of the WVS traditional values index (TRADRAT5) from Wave 6 of the survey. This exercise demonstrates an important feature of MrsP relative to classical MrP; because synthetic poststratification requires only marginal distributions for the weights matrix, we can include demographic information from multiple sources. In this case, because the US Census does not include information on religious affiliation, I could not include it in a classical MrP estimate. With synthetic poststratification, I can use non-Census information, like the Pew Religious Landscape Survey, to fill in that demographic information.

In the first stage, I model TRADRAT5 as a linear function of six demographic variables – age, sex, education level, ethnicity, marital status, and religion (Christian, Non-Christian, or No Religion) – and Census region. A select set of coefficients from this regression is reported in Table 4.

Table 3: Select regression coefficients, regressing TRADRAT5 on six demographic variables and Census region.

	Dependent variable:
	TRADRAT5
White Female	−0.029
	(0.042)
Black Male	−0.412***
	(0.094)
Age 65 -	−0.234***
	(0.067)
University Graduate	0.466***
	(0.165)
Married/Widowed	−0.320***
	(0.043)
Religion: Non-Christian	−0.116
	(0.098)
Religion: Christian	−0.885***
	(0.041)
Region: East South Central	−0.472***
	(0.115)
Observations	2,003
R^2	0.324

I then weight this model's predictions using a synthetic poststratification matrix, with data drawn from the 2010 US Census and the 2014 Pew Religious Landscape Survey. In Figure

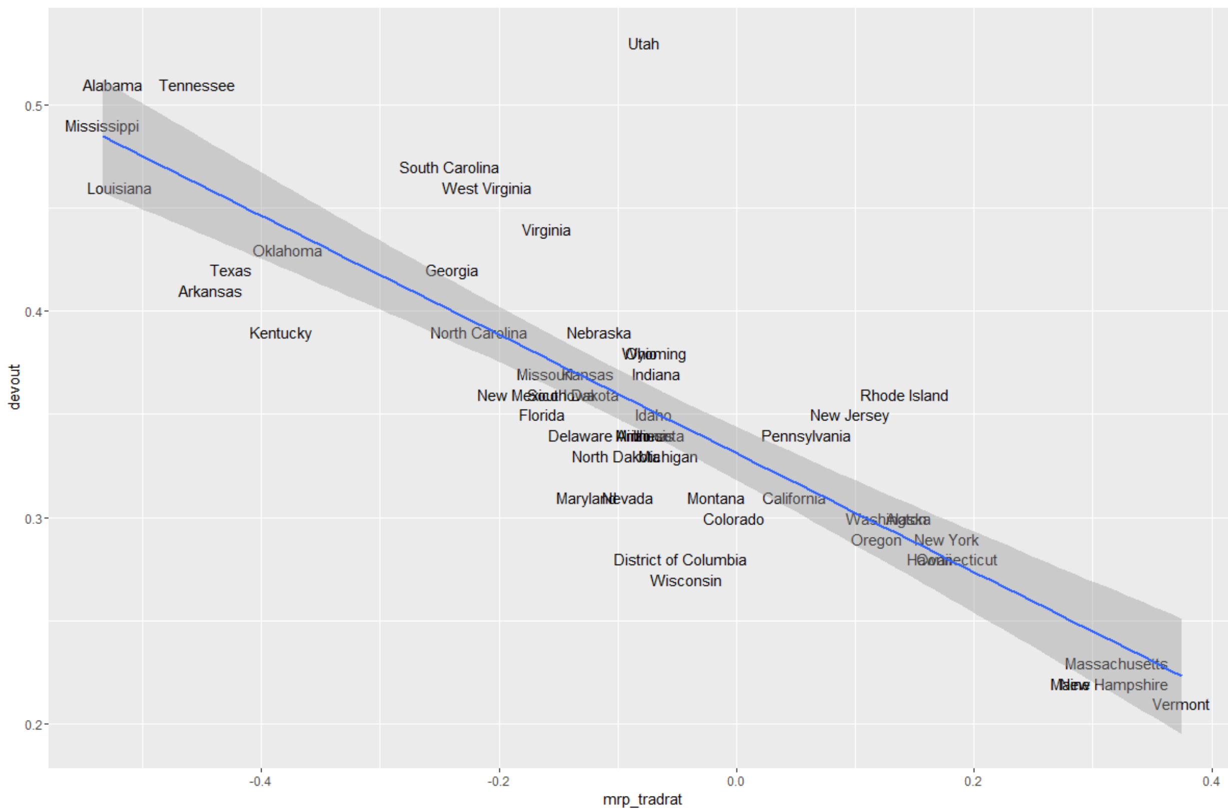

Figure 2: MrsP TRADRAT5 and percentage of Pew survey respondents in each state who attend religious services at least once a week.

2, I compare these estimates against the percentage of adults in each state that report going to religious services at least weekly. As with the Brexit results, the correlation between the MrsP estimates and this alternative measure is high. There is a notable outlier here as well. Utah residents report routine churchgoing more than we would expect given their demographics. No doubt this is due to the heavy concentration of Mormonism in the state, a religious denomination that is not reported in the World Values Survey.

5 Election Forensics (South Africa 2014)

The budding field of election forensics has demonstrated a powerful ability to identify the mathematical signatures of election fraud (Mebane 2008, Klimek et al. 2012). However, while tools like second-digit tests can detect the presence of fraud, it can be difficult to determine from these techniques the extent to which fraud affected the election's outcome. To do so, one would need a "true" measure of support for the winning party, which could then be compared to the reported results. MrsP promises to be a useful complement to the existing election forensics toolkit. By estimating average public support for a political party at the local level, we can judge how well or poorly that party performed relative to its fundamentals. A party that systematically overperforms its fundamentals may just have had a good year, but if other election forensics tools detect anomalous patterns, then MrsP can give us a sense of the size and magnitude of the fraud.

Because MrsP does not rely on detailed joint distributions from Census data, it can be deployed in countries where we were otherwise unable to use classical MrP. For example, South Africa has excellent public census data at the level of its roughly 2,000 electoral wards, but, citing anonymity concerns, does not provide joint distributions for more than two variables at that level. Given this limitation, I use synthetic poststratification to generate ward-level estimates of support for the African National Congress Party (ANC) to compare against observed results from the 2014 provincial elections.

For my survey dataset, I once again turn to the WVS Wave 6. In South Africa, this survey was conducted in 2013, prior to the provincial elections, and asked respondents which party they were likely to support.[1] For the first stage, I estimate a linear-probability model, regressing a binary variable representing the respondent's intention to vote for the ANC on three demographic variables – age, race, and language – and a province dummy. Table 5 reports the coefficient estimates from this exercise.

In the second stage, I weight the model's predictions with a synthetic poststratification matrix. Comparing these MrsP estimates against the observed vote shares for ANC and EFF yields Figure 3. The correlation between the two measures is high, but ANC and EFF substantially outperformed their fundamentals in the Western Cape province, the province where Mebane (2015) reports "extreme fraud probabilities". This is far from conclusive; Mebane also reports "extreme fraud probabilities" in the Free State province, where outcomes do *not* diverge systematically from the MrsP fundamentals. It is, however, a result worthy of further study.

6 Concluding Thoughts

In this paper, I demonstrated that synthetic poststratification can expand the scope of MrP without introducing error. I then presented several empirical applications – the 2016 Brexit referendum, traditional vs. secular-rational values in US states, and the 2014 South African provincial elections – that illustrate the promise of this technique.

MrsP offers three advantages for research in comparative politics. First, it allows researchers to produce public opinion estimates in for local areas that lack census crosstabulations or microdata. This not only allows us to work in countries we couldn't before, but allows us to work in geographic areas (like cities and election precincts) where even in developed countries there are no publicly-available microdata. Second, MrsP permits researchers to generate poststratification weights by combining marginal distributions from multiple datasets. This was particularly useful in the WVS application, given the lack of religious variables in US Census data. Finally, because synthetic poststratification relaxes the data requirements

[1] The survey was also conducted prior to the formation of the Economic Freedom Fighters (EFF) party, a political party that splintered from the ANC. In the validation test, I compare MrsP estimates with the ward-level vote shares for ANC + EFF.

Table 4: Linear-probability model estimates. Dependent variable: intention to vote for African National Congress Party. Selected coefficients on independent variables included here.

	Dependent variable:
	vote.anc
Age 45 - 64	0.083***
	(0.019)
Age 65 -	0.116***
	(0.034)
Race: White	−0.564***
	(0.062)
Race: Other	−0.330***
	(0.070)
Language: Afrikaans	−0.011
	(0.033)
Language: Zulu	0.160**
	(0.066)
Province: Eastern Cape	0.110***
	(0.032)
Province: Limpopo	0.237***
	(0.049)
Observations	3,432
R^2	0.333

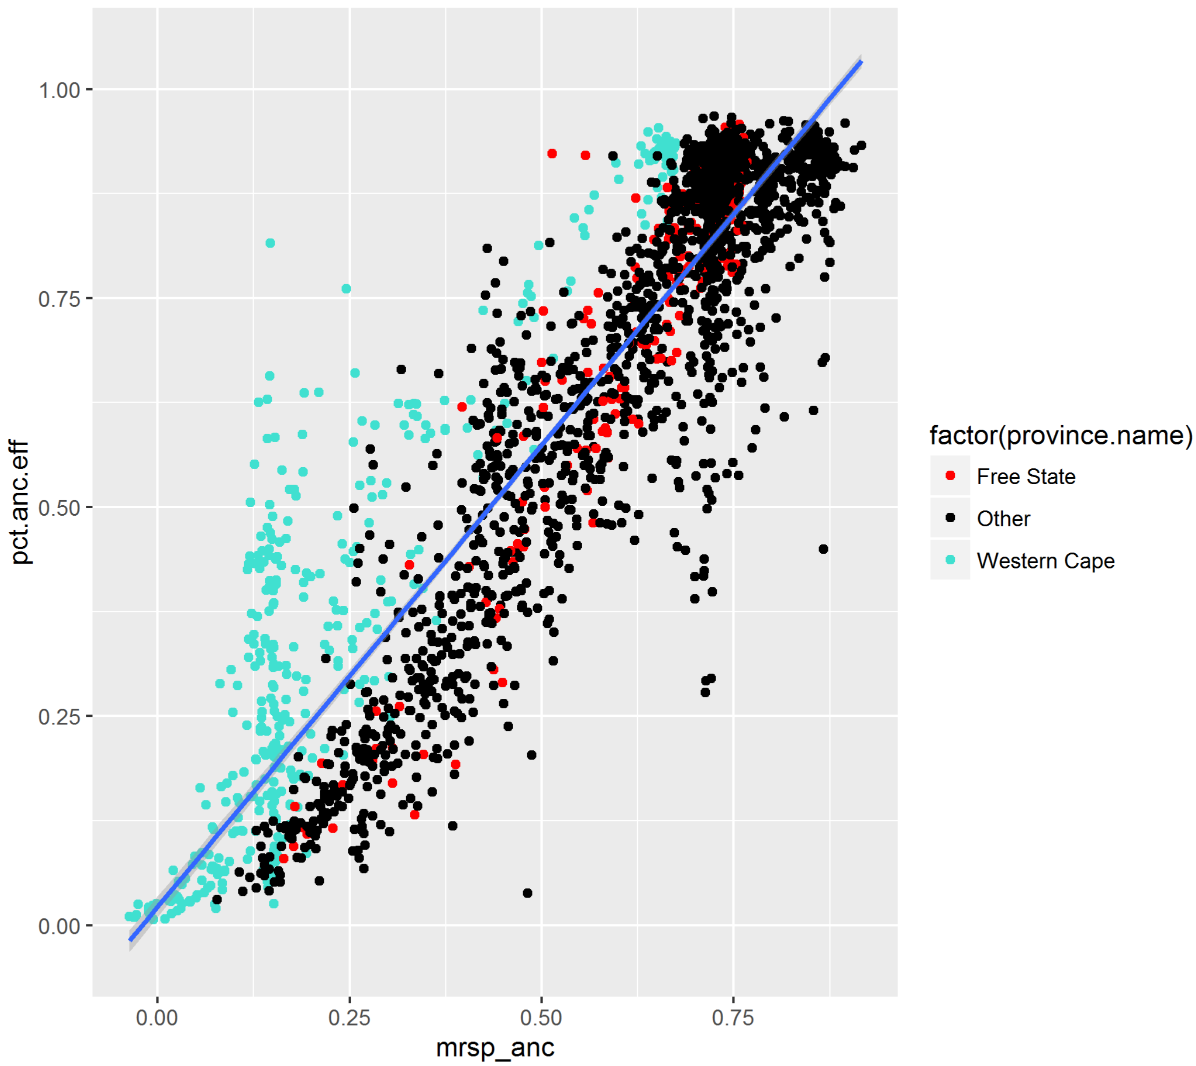

Figure 3: Predicted ANC Support (MrsP) by Electoral Ward and ANC + EFF Vote Share in 2014 South African Provincial Elections

of classical MrP, we can include important individual-level predictors in the model that we previously could only incorporate as geographic predictors. This promises to improve the precision of the method, in addition to its scope.

Going forward, I see three promising avenues for future research. First, MrsP opens up an exciting opportunity to develop subnational politics research in countries with sparse local public opinion data. I hope to see many more applications as the technique becomes more widely adopted. Second, it would be interesting to investigate how well MrsP estimates correlate with precinct-level election data in developed democracies. In the South Africa application, it is unclear whether the observed outliers are the product of mis-estimation, strategic voting, or election fraud. An application in an area where are more confident in the integrity of the results (and voter turnout is high) would be an important comparison case. Finally, it remains an open question how well the first-stage model must fit the data be in order to produce good estimates. Since synthetic poststratification free researchers to include additional predictor variables in the model, the question of how many predictor variables to include becomes more pressing. I hope that future research in this area will incorporate cross-validation techniques, to ensure that the first-stage model is well-fit to the topic at hand.

References

Klimek, P., Yegorov, Y., Hanel, R. & Thurner, S. (2012), 'Statistical detection of systematic election irregularities', *Proceedings of the National Academy of Sciences* **109**(41), 16469–16473.

Lax, J. R. & Phillips, J. H. (2009), 'How Should We Estimate Public Opinion in The States?', *American Journal of Political Science* **53**(1), 107–121.

Lax, J. R. & Phillips, J. H. (2012), 'The Democratic Deficit in the States', *American Journal of Political Science* **56**(1), 148–166.

Leemann, L. & Wasserfallen, F. (2016), 'Extending the Use and Prediction Precision of Subnational Public Opinion Estimation', *American Journal of Political Science* .

Mebane, W. R. (2008), 'Election Forensics: The Second-digit Benford's Law Test and Recent American Presidential Elections', *Election fraud: detecting and deterring electoral manipulation* pp. 162–81.

Mebane, W. R. (2015), 'Election Forensics Toolkit', *DRG Center Working Paper* .

Park, D. K., Gelman, A. & Bafumi, J. (2004), 'Bayesian multilevel estimation with poststratification: State-level estimates from national polls', *Political Analysis* **12**(4), 375–385.

Tausanovitch, C. & Warshaw, C. (2014), 'Representation in Municipal Government', *The American Political Science Review* **108**(03), 605–641.

WORLD VALUES SURVEY Wave 6 2010-2014 OFFICIAL AGGREGATE v.20150418. World Values Survey Association (www.worldvaluessurvey.org). Aggregate File Producer: Asep/JDS, Madrid SPAIN (n.d.).