The Big, The Bad, and The Ugly: Geographic estimation with flawed psychological data

Joe Hoover

Departments of Psychology

Morteza Dehghani

Department of Psychology and Computer Science

University of Southern California

Los Angeles, CA 90089 USA

Author Note

Correspondence regarding this article should be addressed to Joe Hoover, jeh Hoover@usc.edu, 362 S. McClintock Ave, Los Angeles, CA 90089-161. This research was sponsored in part by the Army Research Lab. The content of this publication does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

Many thanks to Lucas Leeman, Håvard Rue, Finn Lindgren, and Andrew Gelman for their thoughtful input on some of the conceptual and computational challenges faced during this work. We also thank the Project Implicit for providing such an accessible and organized public data-base.

Abstract

The geographic distribution of psychological constructs has attracted increasing interest among psychological researchers. In part, this interest has been motivated by the availability of large-scale data collected via online sampling mechanisms. Relying on this and other data, psychologists have been able to not only address novel questions about the spatial variation of psychological constructs but also investigate the regional association between psychological phenomena and real-world outcomes, such as outcomes associated with health, pro-social behavior, and racial inequity. Unfortunately, there many challenges to estimating a construct’s regional distribution — so called *sub-national estimation* — and these challenges are exacerbated by issues of non-representativeness and geographic sparsity. In this work, we provide a comprehensive discussion of major obstacles for small-area estimation and introduce readers to a range of state-of-the-art approaches that rely on multilevel regression and post-stratification (MrP) to deal with these obstacles. We also present a novel evaluation of MrP and extensions of MrP under varying conditions of sample size and response bias (Study 1) and sparsity (Study 2). Finally, in Study 3, we investigate the disagreement between county-level estimates of implicit racial bias obtained via a range of different approaches to sub-national estimation, which we apply to 690,552 responses to the Race IAT collected by Project Implicit. Beyond offering a comprehensive introduction to state-of-the-art approaches to sub-national estimation, this work provides strong evidence for the necessity of incorporating more sophisticated approaches to sub-national estimation into studies of the geographic distribution of psychological phenomena.

Keywords: Sub-national estimation, geographic psychology, multilevel regression and post-stratification, response bias, spatial

The Big, The Bad, and The Ugly: Geographic estimation with flawed psychological data

The regional distributions of psychological constructs are attracting increasing interest in the psychological literature. By studying a construct’s regional variation across space, researchers can learn about its stability, relationships with covariates, and responses to naturally occurring perturbations. For example, a growing body of literature has identified systematic sub-national geographic covariance among personality traits (Allik et al., 2009; Rentfrow et al., 2008, 2015) and between personality and other outcomes, such as life-satisfaction (Jokela et al., 2015), liberalism (Rentfrow et al., 2015, 2013), cancer (McCann, 2017b), volunteering (McCann, 2017a), and work satisfaction (McCann, 2018). Recent research has also provided evidence that the congruence between a person’s personality and the dominant personality traits in their region is associated with their subjective well-being (Götz et al., 2018).

Another burgeoning line of work has focused on the geographic distribution of racial bias and its association with indicators of racial inequity. Studies in this area have identified links between regional implicit bias against Blacks and the Black-White infant mortality gap (Orchard & Price, 2017), Black’s death-rates (Leitner et al., 2016a,b), exposure to racial out-groups (Rae et al., 2015), and disproportionate use of lethal force against Blacks in policing (Hehman et al., 2017). While researchers have long speculated that such associations exist, they have remained difficult to assess quantitatively. However, by focusing on regional variation in target outcomes, researchers have been able gain novel insight into the relationships between psychological phenomena and real-world outcomes.

However, these benefits do not come cheaply, as valid spatial inference requires that estimates of spatial parameters, such as regional means for a given construct, are unbiased. When this requirement is not met — such as when regional estimates are derived from data that is not regionally representative — inference can become unreliable and inaccurate.

Thus, under a direct estimation approach, obtaining unbiased estimates of regional

parameters, such as state- or county-level means, for a given construct requires sufficient and representative sample coverage of the target populations within each region. This can induce costly sample size requirements as the number of regions or the dimensionality of the target population demographics increase. Accordingly, there has been considerable focus on developing strategies and methods that obviate the necessity of collecting very large representative samples while still facilitating sub-national estimation (Erikson et al., 1993; Little, 1993; Gelman & Little, 1997; Park et al., 2004; Selb & Munzert, 2011; Leemann & Wasserfallen, 2017; Hanretty et al., 2016).

However, much of this work remains largely unknown and underutilized among psychological scientists. Indeed, the recent uptick in geographical investigations of psychological constructs has largely been facilitated not by methodological innovation, but rather by the availability of large-scale data-sets collected through non-traditional sampling mechanisms, such as Project Implicit, MyPersonality, and YourMorals.org. To an extent, these data-sets — which consist of hundreds of thousands, millions, or more responses — solve the problem of data availability, as they provide coverage of many sub-national units and demographic populations.

Operating on such data, researchers have conducted ground-breaking studies on constructs such as personality (Rentfrow et al., 2013; Rentfrow & Jokela, 2016) and racial bias (Leitner et al., 2016b; Leemann & Wasserfallen, 2017; Hehman et al., 2017). However, while the sheer size of large-scale internet samples can help researchers overcome some of the obstacles occluding geographical psychological research, obstacles that are independent of sample size, such as response biases, require additional treatment. Unfortunately, because these data are largely generated by internet users self-selecting into participation, the risk of response bias is considerable and this raises serious questions about the validity of using them to derive sub-national estimates (Pew Research Center, 2018; Gelman, 2018). For example, if a data-set over-represents certain demographic populations, regardless of the sample size, sub-national estimates will be biased and estimates for regions populated

by *under*-represented demographic populations may be completely invalid.

A second problem, which is related to but distinct from response bias, is geographic sparsity. Even in very large samples, without rigorously controlled sampling procedures, there is considerable risk of collecting insufficient data from some sub-national units. Further, this risk is compounded by demographic response biases, which can lead to completely or nearly missing data for regions primarily inhabited by populations with low response rates. For example, even with millions of responses, the Project Implicit data is missing or nearly missing responses from many counties. To deal with sparsity, researchers often exclude regions with $N \leq c$, where c is some pre-specified threshold. However, when missingness is correlated with a target outcome, this exclusion policy can lead to false inferences.

These issues raise a pressing challenge: new approaches to collection have yielded data that offer a potential opportunity to investigate questions that previously were inaccessible. However, this data is likely biased and standard analyses may lead to erroneous conclusions. How should researchers proceed? In this work, we address this question via three studies that evaluate the performance of state-of-the-art approaches to sub-national estimation. Specifically, we focus on multilevel regression and post-stratification (MrP; Gelman & Little, 1997; Park et al., 2004), spatial smoothing (Beale et al., 2010; Selb & Munzert, 2011), and multilevel regression and synthetic post-stratification (MrsP Leemann & Wasserfallen, 2017), a recent extension to MrP.

MrP is a model-based approach to obtaining sub-national estimates for a given outcome. It has become increasingly popular and is now considered the gold standard for estimating sub-national political preferences (Leemann & Wasserfallen, 2017; Selb & Munzert, 2011). Recent work has also demonstrated that MrP can even generate surprisingly accurate sub-national predictions from non-random and non-representative data (Wang et al., 2015). However, previous comparative evaluations of MrP have found that it offers diminishing returns as sample sizes increase (Lax & Phillips, 2009; Buttice &

Highton, 2013; Hanretty et al., 2016), suggesting that when ample data is available, simpler approaches like disaggregation (Erikson et al., 1993) — merely calculating region-specific sample means — may perform comparably. These evaluations, however, were conducted with randomly sampled, nationally representative data and thus cannot necessarily be generalized to the kinds of non-random, biased, and big data psychological researchers are working with today. Accordingly, in Study 1, we evaluate the performance of MrP under varying conditions of sample size and response bias.

In Study 2, we focus on a modification of MrP that enables it to better deal with geographic sparsity: the inclusion of a conditional autoregressive (CAR Besag et al., 1991) prior in the MrP model. The CAR prior allows the MrP model to account for spatial auto-correlation and, consequently, this induces a kind of local smoothing, where the estimates for a given region are conditioned on the estimates for neighboring regions. Previous research has shown that this approach can improve MrP estimation accuracy; however, this work has also suggested that, conditional on a strong MrP model, the overall benefits of local spatial smoothing will be incremental at best (Selb & Munzert, 2011; Hanretty et al., 2016). In contrast, rather than focusing on whether local spatial smoothing improves accuracy overall, in Study 2, we investigate whether spatial smoothing improves accuracy as a function of regional sample size. Specifically, we expect that the degree to which local smoothing will improve the accuracy of a region’s estimate is inversely related to that region’s sample size. If this is the case, given the conditions of geographic sparsity that psychological researchers often face, spatial smoothing may be a particularly valuable tool for such work.

Finally, in Study 3, we investigate how county-level estimates of regional racial bias, derived from Project Implicit data, change depending on what estimation procedure is employed. This study enabled us to pursue two goals. First, it offers an opportunity to comprehensively demonstrate how the MrP approach can be applied to problems of sub-national estimation in psychological research. Second, it allows us to observe first-hand

how MrP methods perform with data that is very large, but also non-random and non-representative. In addition to MrP and MrP with spatial smoothing, we also evaluate MrsP (Leemann & Wasserfallen, 2017), a recent modification of MrP that increases the set of demographic variables that can be conditioned on in the MrP response model. This advance is particularly important for psychological researchers, because many psychological outcomes have well-known demographic covariates that could not previously be included in MrP analyses. MrsP thus expands the scope of MrP and considerably increases its use-value for psychologists. More generally, our aim for this study is to investigate the degree to which estimates really change when sophisticated methods for sample adjustment are applied to this kind of data.

Together, these studies evaluate the performance of state-of-the-art methods for sub-national estimation and sample adjustment when applied to data that may be large, but is also non-random, non-representative, and non-uniformly sparse. Specifically, we provide evidence, that under conditions of response bias, MrP can dramatically improve the accuracy of sub-national estimates, regardless of sample size. We also show that while spatial smoothing may not dramatically improve overall accuracy, it can have a substantive positive effect on estimates for regions with few observations. Finally, we demonstrate that even with a very large sample, estimates derived via conventional and spatially smoothed MrP and MrsP diverge substantially from estimates derived from more conventional approaches, such as disaggregation and post-stratification. While we also find that *in this study* there is little estimation variance across MrP-based procedures, analyses conducted with different data or different variables would likely not yield such homogeneity. This is because the degree to which MrsP or spatial smoothing affect estimates, relative to MrP, depends on characteristics of the data being used to obtain estimates.

Sub-national Estimation

Overview

Sub-national estimation involves estimating a distribution of parameters for a set of geographic units that fall below the national level, such as states, provinces, counties, or districts. For example, the problem of estimating the distribution of state-level means for Extroversion or implicit racial bias is a problem of sub-national estimation. Notably, sub-national estimation is neither inherently difficult nor complicated by definition. As is the case with many problems of estimation, access to sufficient data renders the problem trivial. Unfortunately, researchers rarely have access to sufficient data; thus, various methods have been developed in order to facilitate the derivation of sub-national estimates from less-than-ideal data.

One simple and widely-used method is disaggregation (Erikson et al., 1993), which involves combining data from multiple nationally representative surveys into a single data-set and then segmenting the data into the desired level of analysis (e.g. states or counties). Disaggregation hinges on the premise that combining multiple random and nationally representative samples will eventually produce a super-sample that is sufficiently representative at the targeted sub-national level.

While disaggregation is asymptotically valid, in many instances it is not a viable option. While it may be possible to construct a sufficient super-sample for a small set of constructs for which data is frequently collected, this is often not the case for constructs excluded from that set, such as personality inventories and measures of explicit and implicit attitudes. Further, depending on the level of geographic analysis, it may not be possible to assemble a super-sample for even the most widely collected variables. Consider, for example, that moving from a US state-level analysis to a county-level analysis increases the number of spatial units by a factor of approximately 60; thus, data deemed sufficient for a disaggregated state-level analysis would need to be expanded by roughly the same factor to provide comparable coverage for a county-level analysis.

An even more pressing problem for disaggregation is its inability to address response biases and failures in randomization. If certain segments of the target population over- or under-respond, disaggregated estimates will be biased (Holt & Smith, 1979; Little, 1993; Lax & Phillips, 2009) even if they are derived from an infinite sample (Pew, 2018).

To address issues of response bias, researchers employ a range of techniques that aim to adjust sample estimates so that they reflect known population characteristics. Generally, this is accomplished by first using a set of auxiliary variables to segment or *post-stratify* respondents into distinct demographic categories for which sample means are obtained. Then sample weights are derived for each category as a function of that category's representation in the sample and its population count or proportion in the real-world, as estimated via a census or some other source of population statistics. Finally, an estimate of a given sub-national unit's mean is obtained by calculating the weighted mean of the post-strata (e.g. each cross-classification of the demographic variables) sample means, where the weights reflect the demographic population proportions corresponding to the post-strata within the sub-national unit.

Take, for example, the goal of estimating $\mu_{u[l]}^Y$, the population average of the response Y in sub-national unit $u[l]$, where u represents an upper-level geographic boundary system and $[l]$ indexes the lower-level units nested in u ¹. Given this goal, age might be chosen as an auxiliary variable. Then, weights for each post-stratum j of *age* are derived for each sub-national unit $u[l]$ as:

$$w_{u[l],j} = \frac{N_{u[l]} P_{u[l],j}}{n_{u[l],j}} \quad (1)$$

where $N_{u[l]}$ is the total sample size in sub-national unit $u[l]$, $P_{u[l],j}$ is the proportion of people of $age = age_j$ who live in sub-national unit $u[l]$, and $n_{u[l],j}$ is the sample size for

¹ For example, the upper-level unit could be country and, for a single country analysis $U = [1]$ and it contains all sub-national units. While this notation may seem superfluous in this example, we adopt it because it can also be equivalently applied to analyses with multiple upper-level units.

people in sub-national unit $u[l]$ who fall in post-stratum unit j . The weights $w_{u[l],j}$ can then be used to obtain the post-stratified estimate $\bar{Y}_{u[l]}^{ps}$ of $\mu_{u[l]}^Y$:

$$\bar{Y}_{u[l]}^{ps} = \frac{1}{N_{u[l]}} \sum_i^I w_{u[l],i} y_{u[l],i} \quad (2)$$

where $w_{u[l],i}$ and $y_{u[l],i}$ are the weights and outcomes, respectively, for participant i in sub-national unit $u[l]$.

Alternatively, post-stratification estimates can also be obtained without prior calculation of weights:

$$\bar{Y}_{u[l]}^{ps} = \sum_j^J P_{u[l],j} \bar{y}_{u[l],j} \quad (3)$$

where $\bar{y}_{u[l],j}$ is the sample mean for post-stratum j in sub-national unit $u[l]$.

Note that under this approach, the post-stratified mean for a given sub-national unit is function of the post-stratum sample means *within that sub-national unit*. Accordingly, this approach minimally requires $n_{u[l],j} > 0$; however, it is generally preferable to have larger sample sizes, such as $n_{u[l],j} \geq 50$, in order to minimize the effects of sampling error.

Of course, adjustment procedures are often far more complex than this single-variable example. For instance, it is often desirable to select multiple auxiliary variables, with age, gender, race, education being the most used set. However, this can dramatically increase the number of demographic cross-classifications, particularly considering that they are crossed with sub-national units. For example, post-stratifying on 3-level age and education and 2-level gender would produce 18 demographic cross-classifications, which themselves are nested in sub-national units. For a state- or county-level analysis, this approach would yield approximately $18 \times 50 = 900$ or $18 \times 3,007 = 54,126$ distinct participant cross-classifications and adhering to a ≥ 50 rule would require sample sizes of approximately 45,000 or 2.7 million, respectively.

To mitigate such exploding sample size requirements, post-stratification can be reformulated so that estimates of the post-stratum means are *pooled* across sub-national units. That is, rather than estimating the mean for each post-stratum *within* each

sub-national unit — the no pooling approach — the post-stratum means can be estimated across all sub-national units (Gelman & Little, 1997). Thus, the means for each cross-classification of the post-stratification variables are estimated using the entire dataset and sub-national estimates for a given region are obtained simply by weighting the post-stratum means by that regions population weights.

A major benefit of the pooled approach is that the sample size n_j for each post-stratum is much larger and estimates are thus more precise. However, this benefit comes at a cost: the pooling assumes that the post-stratum means are constant across sub-national units. For example, the mean for a given cross-classification of age and gender is assumed to be the same across US counties. Accordingly, while un-pooled post-stratification risks high-standard errors and inflated between-unit variation, pooled stratification risks homogeneity and suppressed between-unit variation.

Another issue that often complicates post-stratification is the difficulty of obtaining population estimates for the cross-classification of the auxiliary variables. Because post-stratification requires an estimate of the full joint distribution of the post-strata cross-classifications, it is often not possible to execute the ideal post-stratification. In such cases, researchers must compromise on the post-stratification design and/or level of sub-national analysis. In response to these issues, methods such as raking (Deville et al., 1993; Kalton & Flores-Cervantes, 2003) are often substituted for post-stratification. Whereas post-stratification operates on the joint distribution of the auxiliary variables, raking operates on their marginal distributions, such that sample weights are derived by iteratively adjusting the marginal distributions of the auxiliary *sample* variables to match the population marginal distributions. For example, raking over age and education at the state-level would involve weighting respondents within each state so that the weighted distribution of their ages matches the known marginal state-level distribution of age. The same procedure would then be applied to education and if this re-weighting interferes with the age alignment, it would be reapplied to age. This iterative re-weighting process would

be repeated until the marginal distributions of the auxiliary variables match their known marginal distributions within some a priori range of error.

Accordingly, raking can considerably expand the pool of viable auxiliary variables, compared to post-stratification. However, the fine-grained information encoded by the full joint-distribution is lost and this can negatively impact estimates. To mitigate this loss of information, raking can also be conducted over some subset of the cross-classifications of the auxiliary variables. However, as with post-stratification, this introduces additional data requirements: the distribution for the chosen cross-classifications must be known and sufficient sample data must be available for each cross-classification category, otherwise estimates may be wildly inaccurate (Gelman, 2007)

Accordingly, the choice between post-stratification and raking is often a choice between a more accurate, but narrower representation of the population (i.e. one that spans only a small set of auxiliary variables) and a less accurate, but wider representation of the population (i.e. one that spans a larger set of auxiliary variables). However, for estimates generated by either approach to be accurate, minimally sufficient data is required for each sub-national unit (unless complete pooling is used). Even for state-level analyses, this can be a daunting constraint as more auxiliary variables are introduced to the estimation procedure.

Multilevel Regression and Post-stratification

While post-stratification and raking remain viable approaches to survey weighting, a more recently developed method, multilevel regression and post-stratification (MrP; Park et al., 2004; Gelman & Little, 1997), has become increasingly popular and is now considered the gold standard for estimating sub-national political preferences (Leemann & Wasserfallen, 2017; Selb & Munzert, 2011). For example, it has been used in sub-national studies on legislative responsiveness to constituent opinion (Kastellec et al., 2015; Krimmel et al., 2012), regional variations in environmental opinions (Fowler, 2016; Howe et al.,

2015), and the relationship between income and political preferences (Gelman, 2014). In contrast to conventional post-stratification, in which sample weights are applied directly to the sample means for each post-stratum, MrP involves applying sample weights to estimates of post-stratum means derived from a hierarchical model fit to individual-level data (Park et al., 2004; Lax & Phillips, 2009). Sub-national means are then estimated as the population-weighted mean of these predicted post-stratum means.

The primary advantage conferred by MrP arises from how post-stratum means for a given outcome are estimated. First, individual-level responses are modeled as a hierarchical, or multilevel, function of demographic auxiliary variables, sub-national geographic indicators, and contextual factors (Park et al., 2004; Lax & Phillips, 2009). Then, this model is used to generate predictions for each cross-classification of the auxiliary variables, conditional on sub-national location and contextual factors.

For example, an individual i 's response y_i , could be estimated as a function of their age, level of education, and sub-national unit (SNU):

$$y_i = \beta_0 + \alpha_{a[i]} + \alpha_{e[i]} + \alpha_{c[i]} \quad (4)$$

Where the effects of the auxiliary variables age and education are modeled as random effects (Raudenbush & Bryk, 2002; Steenbergen & Jones, 2002; Gelman & Hill, 2006), such that they are assumed to be generated from a normal distribution with $\mu = 0$ and variance σ^2 :

$$\begin{aligned} \alpha_a &\sim \mathcal{N}(0, \sigma_a^2), \text{ for } a = 1, \dots, A \\ \alpha_e &\sim \mathcal{N}(0, \sigma_e^2), \text{ for } e = 1, \dots, E \end{aligned} \quad (5)$$

This approach to estimating demographic effects helps address the issue of sparsity in the cells of the cross-classifications of the auxiliary variables, as the effects of a given demographic category are subject to the smoothing effects of partial pooling and shrinkage

toward the mean (Gelman & Hill, 2006; Steenbergen & Jones, 2002; Gelman, 2007). Thus, if there are few observations for a given cross-classification, the estimated effect for that cross-classification will shrink toward the mean of the other levels.

Further expanding on the standard post-stratification approach, MrP also allows modeling the effect of sub-national unit as a random effect. This confers two advantages, first the idiosyncratic effects associated with a given sub-national unit are directly modeled, which helps incorporate between-unit variation into the post-stratification estimates. Second, the parameters fit for each sub-national unit are smoothed toward the upper-level units (e.g. state, region, or country) in which they are nested. Accordingly, the parameter fit for a sub-national unit for which little data is observed will shrink toward upper-level estimates, which helps regularize and stabilize post-stratified estimation.

However, rather than just modeling sub-national units as unconditional random effects, MrP also allows researchers to include relevant contextual factors and, often, additional hierarchical layers. Thus, in many MrP models the effect of sub-national unit might be estimated as a function of unit-level factors and the upper-level unit (ULU) that it is nested in:

$$\alpha_c^{SNU} \sim \mathcal{N}(\alpha_{s[c]}^{ULU} + \beta \mathbf{X}_c, \sigma_c^2), \text{ for } j = 1, \dots, J \quad (6)$$

All together, this approach helps address issues driven by data sparsity. Even if there are no observations for a particular county, its effect can still be estimated as a linear combination of demographic effects, its contextual variable scores, and the effects for the other counties in its state.

Finally, the effect of state can be modeled as a normally distributed random, unconditional effect:

$$\alpha_s^{ULU} \sim \mathcal{N}(0, \sigma_{ULU}^2), \text{ for } j = 1, \dots, J \quad (7)$$

Though, of course, upper-level-unit factors or an additional hierarchical layer could be included to estimate the conditional random effect of the upper-level unit.

The next step in MrP is to use the trained model to generate predictions π_{aec} for each cross-classification of the auxiliary variables age and education, conditional on sub-national location and contextual factors. Thus, in the case of this example, for *each* sub-national unit, a prediction is made for each combination of the levels of age and education. That is, the model is used to estimate the average response for a person of age a with level of education e who lives in county c , for all combinations of $a = 1, \dots, A$, $e = 1, \dots, E$, and $c = 1, \dots, C$.

Finally, post-stratification proceeds similar to as discussed above: sub-national means $\bar{Y}_c^{MrP}$ are estimated by summing over the products of the predicted means and population proportions for each cross-classification of age and education:

$$\bar{Y}_c^{MrP} = \sum_a^A \sum_e^E P_{aec} \pi_{aec} \quad (8)$$

Where, again, P_{aec} is the proportion of people in *sub-national unit* = c of *age* = a with *education* = e and π_{aec} is the predicted outcome for the same cross-classified group.²

As in standard post-stratification, a given sub-national mean estimate from MrP reflects an expected outcome, conditional on the population occupying the sub-national unit. However, whereas standard post-stratification obtains these estimates either through the unpooled approach (i.e. by estimating post-stratum means for each county using only the data observed in that county) or pooled approach (i.e. estimate post-stratum means

² For continuity with the previous and next sections, note that EQ 8 can be written equivalently as:

$$\bar{Y}_{u[s[l]]}^{MrP} = \sum_j P_{u[s[l]],j} \pi_{u[s[l]],j} \quad (9)$$

where $u[s[l]]$ indexes county l , nested in state s , which is nested in national unit u and j indexes each post-stratum of age x education.

using all observed data), MrP affords a middle-ground through partial pooling and is consequently more robust to sparsity within the cells of the demographic post-stratification variables. This robustness to sparsity makes it possible to include more relevant auxiliary variables, which can further improve estimates (Gelman, 2007).

Another notable benefit of MrP is that it is easy to expand the predictive model to exploit known or expected effects. For example, interactions between auxiliary variables can also be estimated and/or the effects of auxiliary variables can be permitted to vary across spatial units, such as regions, which would constitute a so-called random slopes model (Raudenbush & Bryk, 2002).

However, while MrP offers notable advantages, evaluations of its performance highlight that designing a MrP model requires careful thought (Lax & Phillips, 2009; Buttice & Highton, 2013; Hanretty et al., 2016), as its capacity to capture regional variation in an outcome depends on the variance that the predictive model explains. Both Buttice & Highton (2013) and Hanretty et al. (2016) emphasize that well-chosen contextual variables are essential for MrP estimation, finding that MrP models with poor or unrelated contextual variables may not perform well. Similarly, post-stratification variables should not be chosen arbitrarily, but rather with attention to the goal of explaining as much variance in the outcome as possible.

Multilevel Regression and Synthetic Post-stratification

While demographic post-stratification variables should be chosen in order to maximize explained variance in the outcome, there is a strong constraint on whether a variable is eligible for being chosen: the joint distribution of the post-stratification variables' cross-classifications must be known. Otherwise, post-stratification cannot be performed. This has been a major obstacle for MrP estimation, because sub-national joint distributions are not available for many combinations of variables (Leemann & Wasserfallen, 2017). This is particularly relevant for much of psychological research, as

many psychological outcomes are known to covary strongly with demographic constructs like political ideology and religiosity, neither of which have known joint distributions with other standard demographic variables.

Traditionally, this limitation has been circumnavigated via procedures like raking (Deville et al., 1993; Kalton & Flores-Cervantes, 2003), which operate on the easier-to-obtain marginal distributions of auxiliary variables. However, Leemann & Wasserfallen (2017) recently introduced a procedure for conducting MrP with an estimated, or as they call it *synthetic*, post-stratification joint distribution, which they derive from marginal distributions. They provide evidence that their method, multilevel regression and synthetic post-stratification (MrsP, Leemann & Wasserfallen, 2017), performs better than raking, and comparably to or better than MrP (Leemann & Wasserfallen, 2017). Further, while the synthetic joint distributions used in MrsP are subject to error and bias under some conditions, the cost of this is counterbalanced by the benefit of being able to post-stratify over a more diverse set of variables. Thus, by expanding the set of available post-stratification variables, MrsP may actually provide better estimates than those available under MrP.

Simple MrsP. To generate synthetic post-stratification joint distributions, Leemann & Wasserfallen (2017) propose two approaches, which they refer to as ‘simple MrsP’ and ‘MrsP with adjusted synthetic joint distributions’. Under simple MrsP, synthetic joint distributions are calculated merely as the product of the post-stratification variables’ marginal distributions. Thus, for a given set of post-stratification variables, the joint distributions for each sub-national unit k is estimated as the product of the post-stratification variables’ marginal distributions within each sub-national unit. For example, for an arbitrary sub-national unit k , the simple synthetic joint distribution of 3-level age and education P_{aek} can be estimated as the products of their marginal distributions, as shown in Table 1.

It is also worth noting that simple MrsP can be used to extend a known joint

Table 1

*Example of simple MrsP synthetic joint
distribution for age and education*

	age = 1	age = 2	age = 3	
edu = 1	0.01	0.05	0.04	0.1
edu = 2	0.08	0.4	0.32	0.8
edu = 3	0.01	0.05	0.04	0.1
	0.1	0.5	0.4	

Note: The marginal probabilities for each level of age and education are shown in the row and column margins, respectively. The simple synthetic joint distribution is shown in the interior cells.

distribution to include an additional variable for which only the marginal distribution is known. For example, the joint distribution of age and gender could be extended to include education via the same procedure.

After generating synthetic joint probabilities via MrsP, the estimation procedure is identical to estimation with MrP. That is, sub-national means are calculated as the population-weighted mean of the model predictions for each post-stratification cross-classification. Accordingly, substituting MrsP estimation for MrP estimation is relatively easy and requires little additional domain expertise.

However, a notable short-coming of simple MrsP estimation is that the estimated joint distribution will only be correct when the auxiliary variables are independent. As

they diverge from independence, the synthetic joint distribution becomes a less accurate (Leemann & Wasserfallen, 2017). Given that complete independence is rarely observed, this means that simple synthetic joint estimates will almost always be wrong. However, Leemann & Wasserfallen (2017) find that errors in the synthetic joint distribution do not necessarily induce errors in post-stratified, sub-national estimates.

Specifically, sub-national estimates remain constant, regardless of the synthetic joint distribution, as long as the auxiliary variables are modeled with constant marginal effects. Under these conditions, MrP and simple MrsP yield identical sub-national estimates, even if the synthetic joint distribution estimated in MrsP differs substantially from the true joint distribution used in MrP. However, a biased synthetic joint distribution will yield divergent MrsP estimates when the condition of constant marginal effects is violated.

Thus, for instance, Leemann & Wasserfallen (2017) note that MrsP estimates derived from a probit or logistic response model — which both have non-constant marginal effects — will diverge from MrP estimates. However, even with complete dependence among auxiliary variables, the magnitude of this divergence is generally negligible. Importantly, divergences can also occur with linear regression response models, but only when the marginal effects of the auxiliary variables are non-constant, which occurs when interaction effects are estimated. In this case, the degree of divergence between MrsP and MrP will vary as a function of the magnitude of the interactions.

Adjusted MrsP. However, this *does not* mean that researchers should avoid estimating interactions among auxiliary variables, as ignoring a meaningful interaction will also inhibit estimation accuracy. Instead, when auxiliary variables are correlated and their marginal effects are non-constant, researchers can either accept that some degree of bias will affect simple MrsP’s sub-national estimates, or they can rely on the second approach to estimating synthetic joint distributions, which employs an adjustment procedure to refine the synthetic joint distribution.

The goal of this adjustment procedure is to encode any available knowledge about the

true joint distribution in the synthetic joint distribution. For example, while the joint distribution of a set of auxiliary variables may not be known at the *county*-level — thus making MrsP necessary for county-level estimation — in many cases it can be estimated at the national or even state level. Such higher-level estimates of the joint distribution can then be used as a baseline or template for estimating the synthetic joint distribution for each sub-national unit.

To generate an adjusted synthetic joint distribution for a given level of sub-national analysis, the following data is requisite:

1. First, the upper-level, auxiliary variable cross-classification population counts $N_{u,j}$ must be gathered or estimated, where j indexes the cross-classifications of the auxiliary variables and u indexes the upper-level units in $u = 1, \dots, U$. In some cases, this joint distribution may be available via census data; however, surveys can also be used to estimate it. For example, a nationally representative survey could be used to estimate the joint distribution for a set of auxiliary variables, such as gender, age, and education. However, the quality of adjusted MrsP estimates will depend on the accuracy of the auxiliary correlations encoded in this data. Accordingly, researchers must be careful in deciding whether a given survey is sufficiently reliable, as using unreliable data to adjust MrsP can produce estimates that are inferior to simple MrsP.
2. Next, at the targeted lower level of analysis, if the cross-classification population proportion of any subset of the auxiliary variables is available, this should be obtained. We represent this subset of the joint distribution as $P_{u[l],k}$, where $u[l]$ indexes the lower-level units $l = 1, \dots, L$ in upper-level unit u and k indexes the cross-classifications of the auxiliary variable subset. For example, the county-level joint distribution of age $\times$ gender $\times$ education could be the desired joint distribution; but perhaps only the county-level joint distribution of age $\times$ gender is available. In this case, the county-level joint distribution of age $\times$ gender $P_{u[l],k}$ — which represents

the proportion population proportion of people who fall into age x gender cross-classification k for county l , which is nested in upper-level unit u — should be obtained. If no joint distribution data is available for the targeted level of analysis, $P_{u[l],k}$ will be a marginal distribution for one of the variables in the desired joint distribution. For example, if the county-level joint distribution of age x gender is not available, $P_{u[l],k}$ can be the marginal distribution of age or gender in unit $u[l]$.

3. Finally, at the targeted sub-national level of analysis, the *marginal* population proportions $P_{u[l],m}$ for each variable m not represented in $P_{u[l],k}$ must be obtained. Thus, if age x gender x education is the desired joint distribution, but only the joint distribution of age x gender is available, the *marginal* population proportions of education must be obtained in order to construct the synthetic joint distribution.

Given these data, several things can be treated as known. First, the overall or general relationship among the auxiliary variables is encoded in $N_{u,j}$. Second, if $P_{u[l],k}$ is available, the sub-national joint distribution for some subset of the auxiliary variables is known. Finally, for those variables not included in subset in $P_{u[l],k}$, the distributions $P_{u[l],m}$ tell us their marginal population proportions at the targeted sub-national level. The purpose of adjusted MrsP is then to generate a sub-national synthetic joint distribution $P_{u[l],j}^{MrsPA}$ that accounts for the correlational information encoded in $N_{u,j}$.

To accomplish this, $P_{u[l],m}$ is first used to adjust the marginal distribution of m , the variable being added to the joint distribution, in $N_{u,j}$ for each unit $u[l]$ so that it matches $P_{u[l],m}$. In our ongoing example, this means that the marginal distribution of education in $N_{u,j}$ is adjusted to match the known marginal distribution of education in unit $u[l]$, $P_{u[l],m}$. This adjustment is accomplished by transforming $P_{u,j}$ with a correction factor:

$$cf_{u[l]m[i]} = \frac{P_{u[l]m[i]}}{P_{um[i]}} \quad (10)$$

where $P_{u[l],m[i]}$ is the true marginal population proportion of people that fall into level i of

variable m within sub-national unit $u[l]$ and $P_{u,m[i]}$ is an estimate of the same marginal proportion derived from $N_{u,j}$. For example, if $m = \text{education}$, $P_{u,m[i]}$ is the proportion of people with education level i observed in the national-level data; and $P_{u[l],m[i]}$ is the proportion of people with that education level in sub-national unit $u[l]$. Thus, $cf_{u[l],m[i]}$ is simply the ratio of the true proportion of people with $\text{education} = i$ to the proportion estimated from the national data.

$N_{u,j}$ is then transformed as follows:

$$N_{u[l],k,m[i]}^{adj} = N_{u,k,m[i]} \times cf_{u[l],m[i]} \quad (11)$$

where $cf_{u[l],m[i]}$ is the correction factor for level i of variable m , $N_{u,k,m[i]}$ is the set of cross-classification population counts in unit $u[l]$ for which $m = i$, and $N_{u[l],k,m[i]}^{adj}$ is the cross-classification population counts for unit $u[l]$ that have been adjusted so that the margin of $m[i]$ is the same as the observed margin in $P_{u[l],m[i]}$. This means that, in our example, $N_{u[l],k,m[i]}^{adj}$ is the adjusted population count for the cross-classification of age x gender x education in unit $u[l]$; the adjustment ensures that the marginal distribution of education matches the known marginal distribution of education in county $u[l]$.

Finally, the adjusted synthetic joint distribution is generated by using $N_{u[l],k,m[i]}^{adj}$ to extend $P_{u[l],k}$, the known cross-classification population proportions:

$$P_{u[l],j}^{MrsPA} = P_{u[l],k} \times \frac{N_{u[l],k,m[i]}^{adj}}{\sum_{i=1}^I N_{u[l],k,m[i]}^{adj}} \quad (12)$$

where the second right-hand term is the relative weight of m 's levels for each cross-classification of the k , the subset of auxiliary variables for which the joint distribution is known. Specifically, the numerator is simply the adjusted population count of people who fall into cross-classification j (i.e. cross-classification $k, m = i$) in sub-national unit $u[l]$; and the denominator is the adjusted population count of people who fall into cross-classification j , summed across each level of m . This yields $P_{u[l],j}^{MrsPA}$, the estimated proportion of people in sub-national unit $u[l]$ who fall into cross-classification j .

For example, the numerator of the second term in EQ 12 represents the adjusted population count for the cross-classification gender x age x education. The denominator, on the other hand, represents the adjusted population count of the cross-classification of gender x age summed across each level of education. $P_{u[l],k}$ represents the proportion of people in unit $u[l]$ who fall into cross-classification k of age x gender. Finally, $P_{u[l],j}^{MrsPA}$ represents the estimated proportion of people in unit $u[l]$ who fall into cross-classification j of age x gender x education. Accordingly, in contrast to simple MrsP estimates, adjusted MrsP estimates are enhanced by information about the correlational relationship between the auxiliary variables. To generate estimates for an outcome Y in unit $u[l]$, $P_{u[l],j}^{MrsPA}$ can be substituted for $P_u[s[c]], j$ in EQ 8.

This procedure can be repeated multiple times to further extend $P_{u[l],j}^{MrsPA}$. For example, it could be applied to the adjusted synthetic joint distribution of age x gender x education and the marginal distribution of another variable in order to produce a 4-dimensional joint distribution. However, in some instances there may be no available data on the marginal distribution for a desired variable. In such cases, it is still possible to estimate a synthetic joint distribution. However, rather than relying on marginal distributions to implement adjustments, the marginal distribution for the target variable can be estimated via the full adjusted MrsP procedure. That is, a multinomial response model estimating the proportions of the target model can be used to make predictions for each level of the variable and then these predictions can be post-stratified, yielding an estimated marginal distribution for the target variable (Kastellec et al., 2015; Claassen & Traunmüller, 2018). Then, adjusted MrsP can proceed as above. However, while this approach is feasible, its benefit should be weighted against its cost: if a reasonable response model cannot be estimated, then the predicted marginal distribution will be inaccurate and this will negatively effect the accuracy of the final estimation procedure. Accordingly, researchers should carefully evaluate the importance of including a particular variable, keeping in mind that estimation of that variable's marginal distribution may be biased.

Spatial smoothing

MrP and both versions of MrsP help mitigate the consequences of response bias and geographic sparsity. Largely, they accomplish this through their capacity to efficiently exploit information relevant to the real-world distribution of the outcome. In attempts to improve on MrP, various approaches enhancing this capacity have been proposed. One such proposal focuses on expanding the predictive model used in Mr(s)P so that *spatial* information is incorporated into the post-stratum predictions (Selb & Munzert, 2011; Hanretty et al., 2016).

The benefit of this approach is that the estimated effect for a given sub-national unit is partially pooled with the estimated effects for the sub-national units that it is close to. To see the value of this approach, consider that standard Mr(s)P implementations model the effects of sub-national units as a random effect, where a given unit’s effect is shrunk toward the upper-level effect for its group as its sample size decreases (Park et al., 2004; Gelman & Little, 1997; Gelman, 2007). While this is in many ways advantageous, in some circumstances it can also be problematic. For instance, if a sub-national unit differs considerably from its group, but few observations are collected from that unit, its estimated effect will be over-regularized. Incorporating spatial information into a hierarchical model can help counter-balance this kind of excessive regularization.

To do this, researchers have expanded Mr(s)P models to include spatial random effects (Selb & Munzert, 2011; Hanretty et al., 2016). Like standard hierarchical random effects — or *unstructured* random effects, as they are sometimes referred to in the spatial literature — spatial random effects induce partial pooling over a conditional dependency structure. However, whereas unstructured random effects depend on a hierarchical nesting structure, spatial random effects depend on spatial proximity or adjacency. Accordingly, the spatial random effect for a given sub-national unit is partially pooled with the effects estimated for its neighbors. As such, spatial random effects are able to soak up *spatially structured* residual variance among sub-national units. Thus, continuing the example above,

the estimated effect for the unit would still shrink toward the upper-level parameter, but it would also be partially pooled with the estimated effects for its neighbors. In this sense, localized smoothing can help preserve meaningful heterogeneity among sub-national units.

There are many approaches to estimating spatial random effects (Blangiardo & Cameletti, 2015a; Beale et al., 2010; Banerjee et al., 2014) and entire fields of statistical research have been dedicated to developing and refining spatial models. However, small-area estimation research has primarily relied on conditional autoregressive models (CAR Besag et al., 1991; Besag & Kooperberg, 1995; Selb & Munzert, 2011; Hanretty et al., 2016), which rely on a Markov random field specification. This class of priors encode the Markovian assumption that the spatial random effect $\phi_{u[l]}$ for sub-national unit $u[l]$ is conditionally independent of the spatial effects of all other sub-national units $\phi_{u[-l]}$ (where $\phi_{u[-l]} = \{\phi_{u[1]}, \dots, \phi_{u[l-1]}, \phi_{u[l+1]}, \dots, \phi_{u[L]}\}$ for all u), given $u[l]$, where $u[l]$ represents the neighbors j of sub-national unit $u[l]$. As such, CAR models encode conditional dependence among neighboring units. Thus, in a county-level model, a given county's spatial effect will be estimated as a function of the spatial effects of its neighbors. In the same way that an unstructured hierarchical random effect might pool the county's estimate toward the state mean, a CAR spatial effect pools the county's estimate toward its neighbors.

To implement a CAR model, neighborhoods or adjacencies among units must be defined a priori. Often, adjacency is defined as one-step proximity: if unit $u[l]$ shares a boarder with unit $u[j]$, they are considered neighbors. However, adjacency can be defined using other distance functions, such as one- and two-step proximity or some criterion based on a continuous measure of distance. Regardless of how it is defined, these adjacencies must be represented in a symmetric adjacency matrix $\mathbf{W}$ where each $w_{u[l],u[j]}$ entry represents the adjacency score for units $u[l]$ and $u[j]$, which is often simply encoded using a binary indicator, where $w_{u[l],u[j]} = 1$ if units $u[l]$ and $u[j]$ are adjacent and 0 otherwise. In the remainder of this work, as is common, we assume adjacency is defined as one-step proximity and represented using a binary indicator.

Once an adjacency matrix is determined, a standard Mr(s)P response model can be extended by adding a spatially correlated random effect ϕ :

$$y_i = \beta_0 + \alpha_{a[i]}^{age} + \alpha_{e[i]}^{edu} + \alpha_{c[i]}^{county} + \phi_{c[i]} \quad (13)$$

where $\phi_{c[i]}$ is a term that accounts for the spatial effects of unit c 's neighbors.

While there are several distributional forms that can be used to operationalize ϕ , previous work has mainly relied on the intrinsic conditional autoregressive (IAR Besag et al., 1991; Besag & Kooperberg, 1995) prior:

$$\phi_c | \phi_j \sim \mathcal{N}\left(\frac{\sum_j w_{c \sim j} \phi_j}{d_c}, \frac{\sigma_\phi^2}{d_c}\right) \quad (14)$$

where ϕ_c is the random spatial effect for unit c , ϕ_j is the random spatial effect for unit j for all $j \neq c$, $w_{c \sim j}$ are the entries in $\mathbf{W}$ that equal 1 (e.g. that identify the neighbors of c), d_c is c 's number of neighbors, and σ_ϕ^2 is a global variance term for ϕ . Under this specification, ϕ_c is estimated as the average value of ϕ among c 's neighbors with variance scaled by the number of neighbors that c has. Accordingly, as d_c increases and/or σ_ϕ^2 decreases, the distribution of ϕ_c becomes more tightly clustered around its mean, which is to say more local smoothing is induced.

When an IAR random effect is added to a standard independent and identically distributed (i.i.d.) hierarchical model, the model is conventionally referred to as a Besag-York-Mollié model (BYM Besag et al., 1991; Blangiardo & Cameletti, 2015a). In the BYM model, partial pooling is conducted within the specified hierarchical structure *and* among neighboring units, which leads to a combination of regional and local smoothing (Selb & Munzert, 2011; Hanretty et al., 2016).

The benefit of the BYM model's dual approach to modeling the random effects of sub-national units is that two types of structured residual variance are accounted for. On one hand, the unstructured random effect captures residual variance among sub-national units that are nested in the the same upper-level unit (e.g. region or state); this is a simple

and powerful way to induce pooling over large regions that may share similarities. On the other hand, the spatially structured random effect captures residual variance between neighboring units. This enables the model to be sensitive to spatially structured heterogeneity *within* upper-level units. Further, it allows for residual variance between units that are adjacent but occupy different upper-level units. This second characteristic is important: it accounts for the very real possibility that, for example, if two counties are adjacent, the fact that they are located in different states does not mean they are identical and independently distributed, conditional on their state, which is the assumption made by unstructured random effects models. This all can benefit Mr(s)P estimation, because it improves the exploitation of the information encoded in the observed data.

However, one drawback of the BYM model is that only the sum of the unstructured and spatially structured random effects is identifiable (Lee, 2011; Eberly & Carlin, 2000) and this destabilizes estimation for the individual components (Eberly & Carlin, 2000). This can be particularly problematic when there is little or no residual spatial auto-correlation, because inflated spatial auto-correlation may be falsely detected. To address this, as well as other issues (such as the difficulty of setting reasonable priors), a new parameterization of the BYM model, the modified BYM, was recently proposed (Riebler et al., 2016; Simpson et al., 2017). This new parameterization introduces a mixing parameter that controls the relative contributions of the unstructured and spatially structured residual variance for geographic units. Consequently, the modified BYM is better able to adjust to various combinations of unstructured and spatially structured covariance, such that when no spatial auto-correlation is present, it simply reduces to an unstructured model and if no unstructured residual variance is present, it reduces to an IAR model (Riebler et al., 2016).

Regardless of the functional form used, including a spatial error term in Mr(s)P models *can* improve model performance. However, the magnitude of this improvement depends on several factors. Most importantly, if no spatially structured residual variance is

present in the data, conditional on the non-spatial model, accounting for spatial auto-correlation confers no advantage and may actually impede accuracy, depending on what spatial prior is used. That is, while spatial auto-correlation may be observable in the data, the benefit of the spatial error term depends on whether this spatial auto-correlation persists *after* accounting for the other components in the Mr(s)P response model.

Consistent with this, Hanretty et al. (2016) find that when strong contextual factors are included in MrP models, spatial smoothing may offer little advantage, as the contextual factors explain the covariance between adjacent sub-national.

However, evaluations of when and to what degree spatial smoothing benefits Mr(s)P estimation have focused on global improvements. Accordingly, it is not entirely surprising that, conditional on a strong unstructured Mr(s)P model, spatial smoothing has been found to offer incremental improvements, at best. For units with a reasonable number of observations, spatial smoothing will be minimal except under conditions of high spatial auto-correlation, which are unlikely given a strong unstructured model. Thus, in analyses where most units have a reasonable number of observations, spatial smoothing will have little influence over post-stratum predictions *on average*. However, for the units with few or no observations, this will not be the case, as more spatial pooling will occur, relative to the degree of spatial auto-correlation.

Accordingly, in our view, spatial smoothing is most likely to improve Mr(s)P estimates for sub-national units with few observations, because estimation for those units will be pooled with estimates from adjacent units. To the extent that the observed people who live in sub-national units $u[l]$ and $u[l, j]$ share un-modeled, idiosyncratic similarities, spatial smoothing functionally increases the sample size for $u[l]$. If $u[l]$ already has a reasonably large sample size, this will have little effect on post-stratum predictions; however, when $u[l]$ has few or no observations, anchoring on the data observed in adjacent units should improve model estimates. In Study 2, we investigate this hypothesis by evaluating the accuracy of MrP estimates derived from unstructured and spatially

structured response models as a function of the unit-level observed sample size.

Study 1

Previous evaluations of MrP have found that it offers diminishing returns, compared to disaggregation, as sample sizes increase. For instance, Lax & Phillips (2009) found that disaggregation and MrP performed comparably with a sample size of approximately 14,000 in a state-level analysis with 49 states. However, this convergence in performance is contingent on the sample being representative, because, as sample size is increased, a representative sample systematically approaches the population.

In contrast, when a sample is not representative (e.g. due response biases), increasing the sample size does not necessarily yield a more precise approximation of the population. For example, if a particular population segment is not represented in a sample, estimates for sub-national units populated by that segment may be biased and the degree of that bias will partly depend on the population portion of that segment. Under such conditions, the performance of disaggregation and MrP will still move toward convergence as sample sizes increase as, after-all, an exhaustive sample would be equivalent to the population. However, this convergence will be inhibited by the degree to which the sample departs from representativeness.

To better understand this process, in this study, we investigate the results of Monte Carlo simulations (Mooney, 1997) in which disaggregation and MrP estimates are compared under different conditions of sample size and response bias.

Method

Our Monte Carlo simulation is structured as follows. First, a population of respondents is generated over a grid of 400 sub-national units. Specifically, over the grid of units, marginal distributions of demographic characteristics are sampled for three demographic variables: one two-level, γ^1 variable and two three-level variables, γ^2 and γ^3 . Then, constrained by these marginal distributions, demographic characteristics are

assigned to simulated respondents within each sub-national unit. Thus, each respondent is associated with a specific level of each γ^d , for d in $d = 1, 2, 3$.

Then, given the generated population, a set of population weights are randomly drawn. These weights consist of linear effects for the demographic and contextual factors, a random sub-national unit effect, and individual-level error. These population weights are used to generate values for the response variable $Y_{l[i]}$.

Next, samples of sizes $S = [100, 1,000, 5,000, 10,000]$ are drawn from the population for each of three degrees of response bias. To simulate response bias, respondents are drawn with specified probabilities for each level of one of the three-level demographic variables. Specifically, three different degrees of bias are examined: $p = (1/3, 1/3, 1/3)$, $p = (0.4, 0.35, 0.25)$, and $p = (0.70, 0.2, 0.1)$. Further, to simulate random residual response-bias at the sub-national unit level, response probabilities are randomly assigned to each sub-national unit, which makes responses from some units more likely than others. Finally, given a drawn sample, disaggregation and MrP estimates of unit means are obtained and root mean squared error and bias are calculated.

For this study, three populations ($Ns = 10,000,000$) were simulated. Then, for each population, 50 sets of population weights were sampled. Finally, for each combination of the four sample sizes and three levels of response bias, 100 samples were drawn. This yielded a total of 180,000 iterations, or 15,000 iterations for each combination of sample size and response bias. These settings were chosen in order to minimize uncertainty while also maintaining reasonable computational cost. For a detailed description of the data generating process, please see Study 1 Data Generating Process in Supplemental Material.

Results

For each sample drawn within our simulation framework, disaggregation and MrP estimates were obtained and used to calculate RMSE and bias. To evaluate the performance of these methods as a function sample size and response bias, we estimated

the mean RMSE and bias for each method within each combination of sample size and response bias.

As expected, per previous findings, under conditions of low response bias, the performance of disaggregation quickly converges with that of MrP (See Table 2 and Figure 1). For example, with samples of 100, disaggregation's expected RMSE ($\bar{X} = 0.59, \hat{\sigma} = 0.04$) was 0.22 higher than MrP's ($\bar{X} = 0.37, \hat{\sigma} = 0.04$). However, increasing the sample size to only 1,000 considerably reduced this gap, such that disaggregation's expected RMSE ($\bar{X} = 0.18, \hat{\sigma} = 0.01$) was 0.05 higher than MrP's ($\bar{X} = 0.13, \hat{\sigma} = 0.01$). Further, with samples of 10,000, disaggregation's expected RMSE ($\bar{X} = 0.06, \hat{\sigma} < 0.01$) was only about 0.015 higher than MrP's ($\bar{X} = 0.04, \hat{\sigma} < 0.01$).

Figure 1. RMSE (A) and mean bias (B) as function of sample size and bias. Error bars represent 2.5th and 97.5th percentiles in simulation distribution.

However, while the convergence of disaggregation and MrP's expected RMSE only decreases slightly under medium response bias, under high response bias the convergence is attenuated considerably. Further, our simulations suggest that as response bias increases, disaggregation's RMSE becomes more variable. Notably, this increase in variance is not

observed for MrP.

Regarding the estimation bias of disaggregation and MrP, neither method showed strong mean bias under any conditions (See Table 2 and Figure 1). However, the variances of their estimates of bias demonstrated starkly different patterns. Specifically, while, in the low bias condition, the variance of both disaggregation and MrP's bias shrunk toward zero as sample size increased, this remained true only for MrP as response bias was introduced. That is, as expected, under conditions of response bias, increasing the sample size had virtually no effect on the estimation bias of disaggregation.

Table 2

Mean RMSE and Bias by Sample Size and Bias

Sample Size	Bias	Mean RMSE		Mean Bias	
		Disaggregation	MrP	Disaggregation	MrP
1,000	Low	0.594 (0.04)	0.368 (0.04)	-0.001 (0.03)	0 (0.02)
1,000	Medium	0.595 (0.04)	0.369 (0.04)	0.009 (0.04)	0.001 (0.02)
1,000	High	0.61 (0.04)	0.38 (0.04)	0.022 (0.11)	0 (0.03)
10,000	Low	0.184 (0.01)	0.129 (0.01)	0 (<0.01)	0 (<0.01)
10,000	Medium	0.19 (0.02)	0.13 (0.01)	0.011 (0.04)	0 (<0.01)
10,000	High	0.245 (0.05)	0.14 (0.02)	0.031 (0.13)	0 (<0.01)
50,000	Low	0.08 (<0.01)	0.058 (<0.01)	0 (<0.01)	0 (<0.01)
50,000	Medium	0.091 (0.01)	0.058 (<0.01)	0.011 (0.04)	0 (<0.01)
50,000	High	0.16 (0.06)	0.062 (<0.01)	0.031 (0.13)	0 (<0.01)
1e+05	Low	0.056 (<0.01)	0.041 (<0.01)	0 (<0.01)	0 (<0.01)
1e+05	Medium	0.071 (0.02)	0.041 (<0.01)	0.011 (0.04)	0 (<0.01)
1e+05	High	0.146 (0.07)	0.044 (<0.01)	0.031 (0.13)	0 (<0.01)

Discussion

Previous evaluations of MrP suggested that the performance disaggregation and MrP converges as sample size increases. However, these evaluations were conducted with representative samples. In this study, we show that this convergence is inhibited as response biases are introduced to the sampling mechanism. Specifically, we found that under conditions of response bias, MrP considerably outperformed disaggregation in terms of error and bias, regardless of sample size.

However, it is important to interpret MrP's superior performance in these simulations in context. In any situation, MrP's performance will depend on the association between the outcome and the variables selected for post-stratification, the contextual factor(s), and the modeled residual hierarchical variance (Buttice & Highton, 2013). Accordingly, the degree to which MrP outperforms disaggregation is dependent on strength and comprehensiveness of the MrP model. If a MrP analysis does not include a strong contextual factor or it post-stratifies on demographic variables that explain very little variance in the outcome, the estimates generated by the analysis will be worse than those generated by a stronger MrP analysis. This means that a weak MrP analysis may not substantively outperform disaggregation, as performance depends on the quality of the MrP model. We emphasize this not to suggest that disaggregation is a viable alternative to MrP, but rather to highlight the importance of building a strong MrP model.

Setting aside the issue of building a MrP model, our results suggest that if a sample is possibly subject to response biases, disaggregation should not be used, even if the sample is very large. This is particularly relevant for the analysis of data generated by non-random, voluntary self-reports, because such data is almost certainly subject to a wide range of response biases.

Study 2

The standard hierarchical parameterizations used in Mr(s)P models assume that observations are i.i.d., given the hierarchical structure of the model. In many cases, this assumption may be unreasonable, as within any given region, observations from counties that are closer together may be more similar than those from counties that are further apart. Additionally, observations from counties that are geographically close, but nested in different regions may also exhibit similarities.

The possibility of such spatial auto-correlation suggests that i.i.d. Mr(s)P models may not optimally exploit the information available in the data. At least under conditions of residual spatial auto-correlation, a model that accounts for spatially structured patterns should perform better. However, previous research has found that MrP with spatial smoothing offers, at best, an incremental improvement over i.i.d. MrP (Selb & Munzert, 2011; Hanretty et al., 2016). This most likely a consequence of the fact that, at least for the variables under study, strong MrP models left little residual spatial auto-correlation for spatial smoothing to exploit.

However, it may also be the case that previous work overlooked the situations where spatial smoothing is most likely to improve MrP estimation. Comparable to smoothing in i.i.d. models, the magnitude of spatial smoothing for a given unit (e.g. the degree to which that unit will be influenced by its neighbors) depends on the observed sample size for that unit. Thus, we expect that spatial smoothing should offer the strongest improvements in estimation accuracy for units with missing data or small sample sizes.

If this is the case, spatial smoothing may be particularly beneficial for obtaining sub-national estimates under conditions of sparsity. To investigate this possibility, we conduct a second Monte Carlo simulation and compare estimates obtained via disaggregation, MrP, and MrP with spatial smoothing under varying conditions of spatial auto-correlation.

Method

We use a data generation process that is similar to the procedure developed for Study 1. However, in the current study, we simulate four different degrees of spatial auto-correlation in order to evaluate the effects of spatial smoothing under a range of conditions. We also hold sample size ($N = 1,000$) constant and do not introduce response biases.

Data Generating Process. Similar to Study 1, we simulated two large populations ($N = 10,000,000$) with simulated demographic characteristics over a grid of 400 units. For each population, we again sampled 5 different sets of regression parameters and used these to generate a dependent variable Y . For each instance of Y , we induced one of four degrees of spatial auto-correlation. This was achieved by simulating a spatially structured random effect over the grid of units, which we modeled as a multivariate normal distribution:

$$\rho_j \sim MVN(0, \sigma^2 e^{-\theta d_{jk}}) \quad (15)$$

where σ^2 controls the scale of the spatial random effect, θ controls the decay of spatial auto-correlation, and d_{jk} encodes the distance between units j and k (Guélat & Kéry, 2018). In our simulations, we set $\theta = 0.2$ and adjusted the magnitude of spatial auto correlation via the linear combination:

$$y_{ij} = Y + w_a \times \rho_j \quad (16)$$

where $w_a \in [0, 0.5, 1, 1.5]$ magnifies or reduces the spatial random effect ρ_j .

Finally, for each level spatial auto-correlation, we drew 100 samples with $N = 10,000$, which, collapsing across the 10 total sets of simulated regression parameters, yielded $100 \times 100 = 10,000$ samples drawn for each level of spatial auto-correlation.

Estimation procedure. Aside from the estimation methods, the estimation procedure used in the current study was identical to that used in Study 1. However, here, we generated unit-level estimates from four different methods: (1) disaggregation; (2) MrP with an i.i.d. error structure estimated via maximum likelihood (ML); (3) MrP with an i.i.d. error structured estimated via Integrated Nested Laplace Approximation (INLA) (Blangiardo & Cameletti, 2015a,b; Martino & Rue, 2009), a fast and scalable approximate Bayesian estimation procedure; and (4) MrP with a modified BYM error structure, also estimated via INLA.

By estimating two i.i.d. MrP models, one via ML and one via INLA, we were able to evaluate the effects spatial smoothing on estimation performance while accounting for potential differences caused by changing estimation algorithms.

Results

To evaluate the effects of local spatial smoothing on estimation performance we conducted two analyses. First, we calculated mean RMSE as a function of auto-correlation condition and method of estimation (See Table 3). Results indicated that even under conditions of high spatial auto-correlation, inducing local spatial smoothing did not meaningfully improve overall estimation accuracy. However, consistent with both Study 1 and previous research, we also found that, regardless of estimation method, MrP estimates were more accurate than estimates obtained via disaggregation.

Second, we calculated RMSE as a function of auto-correlation, method, and sample size. Results from this analysis were consistent with our expectation that, under conditions of spatial auto-correlation, local spatial smoothing would yield estimation improvements for units with few observations (See Figure 2). Specifically, for each set of MrP estimates, RMSEs were functionally equivalent under conditions of no and low spatial auto-correlation, regardless of unit sample size. However, under conditions of medium and high spatial auto-correlation, RMSEs for MrP estimates generated from the modified BYM

Table 3

RMSE by spatial auto-correlation and estimation method

Spatial auto-correlation	Estimation method			
	Disaggregation	i.i.d ML	i.i.d. INLA	Modified BYM
None	0.18	0.13	0.13	0.13
Low	0.18	0.13	0.13	0.13
Medium	0.16	0.12	0.12	0.12
High	0.15	0.11	0.11	0.1

All standard errors < 0.001 .

model were consistently lower than those generated from the i.i.d. models (See Table 1 in Supplemental Materials for point estimates). This difference was largest for units with no observations (e.g. units for which a single observation was mean imputed), such that RMSE was, on average, 0.06 points lower for spatially smoothed estimates ($M = 0.55$, $SE = 0.02$) in the high auto-correlation condition and 0.05 points lower ($M = 0.5$, $SE = 0.02$) in the medium auto-correlation condition (See Table 4). However, as unit sample size increases, our results also indicate that the benefits of spatial smoothing rapidly attenuate. For example, for units with a sample size of 5, we observed only a 0.01 decrease in RMSE for spatially smoothed estimates.

Discussion

Previous evaluations of the impact of local spatial smoothing on MrP estimation have indicated that accounting for spatial auto-correlation offers little overall benefit above and beyond a strong i.i.d. MrP model (Selb & Munzert, 2011; Hanretty et al., 2016). However, results from the current study suggest that spatial smoothing may indeed offer substantive

Table 4

RMSE by spatial auto-correlation, method, and sample size

Spatial auto-correlation	n	Disaggregation	i.i.d. ML	i.i.d. INLA	modified BYM
None	0	0.60 (0.02)	0.46 (0.02)	0.46 (0.02)	0.46 (0.02)
Low		0.67 (0.02)	0.53 (0.02)	0.53 (0.02)	0.51 (0.02)
Medium		0.64 (0.02)	0.55 (0.02)	0.55 (0.02)	0.5 (0.02)
High		0.71 (0.02)	0.61 (0.02)	0.61 (0.02)	0.55 (0.02)
None	3	0.36 (<.01)	0.23 (<.01)	0.23 (<.01)	0.23 (<.01)
Low		0.35 (<.01)	0.23 (<.01)	0.23 (<.01)	0.23 (<.01)
Medium		0.33 (<.01)	0.22 (<.01)	0.22 (<.01)	0.21 (<.01)
High		0.29 (<.01)	0.2 (<.01)	0.2 (<.01)	0.2 (<.01)
None	5	0.28 (<.01)	0.19 (<.01)	0.19 (<.01)	0.19 (<.01)
Low		0.27 (<.01)	0.19 (<.01)	0.19 (<.01)	0.19 (<.01)
Medium		0.25 (<.01)	0.17 (<.01)	0.17 (<.01)	0.17 (<.01)
High		0.23 (<.01)	0.16 (<.01)	0.16 (<.01)	0.15 (<.01)

Standard errors shown in parentheses.

gains in estimation accuracy for units with very few observations. For example, under conditions of medium and high spatial auto-correlation, local spatial smoothing induced by a modified BYM prior (Riebler et al., 2016) decreased RMSEs for missing units by 0.05 and 0.06 points, respectively. This suggests that while local spatial smoothing may not have a dramatic impact on overall error, under conditions of sparsity it may substantively improve estimation accuracy.

Importantly, under conditions of no and low spatial auto-correlation, our results also indicate that local spatial smoothing with the modified BYM prior does not decrease

Figure 2. Model RMSE by Spatial Auto-correlation and Cell Sample Size. Disaggregation RMSE is not shown due to consistent and large relative under-performance. M2 = MrP with an i.i.d. error structure estimated with maximum likelihood; M3 = MrP with an i.i.d. error structure estimated with INLA; M4 = MrP with a modified BYM error structure estimated with INLA. Error bars represent bootstrapped confidence intervals.

accuracy. In such instances, estimation accuracy appears to be equivalent to the accuracy of estimates obtained from i.i.d. models. Accordingly, while spatial smoothing may only substantively increase estimation accuracy under certain conditions, our results also suggest that spatially smoothed estimation accuracy should, at worst, reduce to i.i.d. estimation accuracy.

Study 3

In Study 1, we found that while MrP and disaggregation estimates rapidly converge as sample size increases under conditions of no response bias, under conditions of response

bias this convergence rate is much slower. Further, in Study 2, we found that inducing local spatial smoothing, under some conditions, can substantively improve estimation accuracy for units with missing or very few observations. However, both of these studies were conducted in pristine environments and they offer little information about how estimates obtained from various MrP estimation procedures may differ (or fail to differ) from each other and from estimates obtained from more conventional approaches, such as disaggregation and post-stratification, when applied to large data collected via opt-in online survey methodologies.

Accordingly, in the current study, we use disaggregation, post-stratification, MrP, and MrsP to obtain county-level estimates of implicit racial bias using data obtained from Project Implicit (Xu et al., 2013). We include post-stratification as an intermediary estimation procedure because it represents an improvement above and beyond disaggregation, but nonetheless is thought to fall considerably short of MrP due to issues raised by sparsity within the post-stratum cells (Park et al., 2004; Lax & Phillips, 2009; Gelman & Little, 1997).

While we could have selected data from other sources for this study, we chose to focus on Project Implicit’s racial bias data for two reasons. First, it exemplifies the kind of large-scale, high-value data that can be collected via online, opt-in collection strategies. Over 16 years of operation, Project Implicit has collected millions of responses and it offers an unprecedented opportunity to examine geographic variations in racial bias as well as to investigate its associations with other outcomes. Second, because of its value, this data is increasingly being used in geographic studies of racial bias and we expect that such applications will only become more popular in the future. As such, we hope that an accessible exploration of how estimates may vary depending on estimation procedure will be useful for researchers.

To this end, we evaluate variations in estimates obtained from seven different estimation procedures: (1) disaggregation, (2) i.i.d. MrP estimated with ML, (3) i.i.d. MrP

estimated with INLA,(4) MrP with spatial smoothing, (5) i.i.d. MrsP estimated with ML, (6) i.i.d. MrsP estimated with INLA, and (7) MrsP with spatial smoothing. While we do not have a ground-truth measure of implicit racial bias, this approach allows us to investigate whether and to what extent these methods produce divergent estimates. Of course, mere divergence does not provide information regarding accuracy; however, in conjunction with Studies 1 and 2, as well as with previous research, our view is that if divergences occur, estimates obtained from the more sophisticated methods should be preferred as they should only at worst reduce to estimates obtained from simpler methods.

Further, and perhaps more importantly, evaluating divergence can provide information about whether a more sophisticated method is likely to be worth the cost of implementation. For example, inducing local spatial smoothing over a large number of geographic units dramatically increases the computational cost of a model. While our simulation results suggest that the more costly model may substantively improve estimates for units with few observations, it may be the case that when applied to real-world data there is very little divergence between spatially smoothed and i.i.d. estimates. Such an outcome would, of course, be circumstantial, as the influence of spatial smoothing is dependent on the magnitude of residual spatial auto-correlation; nonetheless, observing the estimation variation associated with these methods can help inform researcher’s decisions about how to efficiently minimize computational cost.

Data

Race IAT. The primary data used for this study were responses to the race implicit association test (IAT) obtained from Project Implicit (Xu et al., 2013) collected between 2015 and 2017. The IAT relies on a timed dual-categorization task that requires respondents to evaluate pairings of White and Black faces and words referring to ‘good’ and ‘bad’ things. An indication of racial bias (against Blacks) occurs when a respondent more quickly categorizes words representing "bad" things as bad when they are paired with

a Black face, compared to a White face, and when they are able to more quickly categorize "good" words when they are paired with a White face (Greenwald et al., 2009). Through repeated measures sampling across categorization trials, the IAT permits the estimation of the so called D score, which represents the difference in response latency. D scores range from -2.0 to 2.0, where scores above 0 indicate a positive bias toward White faces and a negative bias toward Black faces.

2011-2015 American Community Survey 5-year estimates. To estimate county-level joint distributions, we rely on US Census data obtained from the 2011-2015 American Community Survey 5-year estimates (Center, n.d.).

Pew 2014 Religious Landscape Study. We also rely on data collected by the Pew organization to extend the joint-distribution used in our MrsP models. Specifically, we use data collected for Pew’s Religious Landscape Study (Center, 2014; Hackett et al., 2014), a nationally representative US data ($N = 35,071$) designed primarily to evaluate trends in American religiosity. We use this data to estimate the marginal distributions of county-level religiosity and political ideology.

Census Bureau’s geographic database. County names and locations were determined using data provided by the US Census Bureau’s geographic database (Bureau, n.d.).

Method

In this study, we estimate county-level Race IAT D scores for the contiguous US using disaggregation, post-stratification by age, and a range of MrP approaches. For our standard MrP estimates, age, gender, and race are selected as post-stratification variables, as these are directly available from the US census. In contrast, for our MrsP estimates, we extend the post-stratification variables to also include level of education and political ideology.

Primary Data preparation. Individual responses to the 2015 ($N = 1,028,518$), 2016 ($N = 1,051,15$), and 2017 ($N = 945,10$) Project Implicit Race IAT data with

complete responses to county location, age, gender, ethnicity, education level, and 7-point political ideology ($N = 168, 231, 284, 558$, and $353, 934$, respectively) were selected for analysis (Total $N = 86, 723$). While complete responses to all of these items are only required for MrsP estimation, this selection process was applied for all methods in order to maintain uniformity. Finally, responses were limited to those associated with the 318 counties located in the contiguous US, $N = 690, 552$. This yielded a dataset with coverage of 2,952 counties (County N summary statistics: *Mean* = 233.9, *Median* = 25.0, 1st quartile = 8.0, 3rd quartile = 10, min = 1, max = 2,352). Finally, race, age, gender, education, and 7-point political ideology were also recoded as follows: race = [Black, Hispanic, Other, White]; age = [2-29, 3-44, 45-64, 65+]; gender = [female, male]; education = [High school graduate or less; some college through bachelors degree; graduate degree³]; ideology = [conservative, moderate, liberal].

Secondary Data preparation. Two contextual variables were selected for the MrP and MrsP models, county-level Democratic vote share for the 2016 Presidential election and a measure of racial segregation, exposure (Massey & Denton, 1988), which is interpreted as the probability that a member of one racial group will meet a member of another target group. Using the standard definition of exposure, we calculated it for Whites to Blacks as:

$$B_{l[c]} = \sum_{i=1}^I \frac{\frac{w_{l[c]}}{W_l}}{\frac{b_{l[c]}}{n_l[c]}} \quad (17)$$

where $w_{l[c]}$ is the number of Whites in census tract c of county l , W_l is the number of Whites in county l , $b_{l[c]}$ is the number of blacks in census tract c in county l , and $n_l[c]$ is the population of census tract c in county l . Accordingly, B_l is interpreted as the weighted exposure of Whites to Blacks in in county l . Notably, exposure does not measure segregation per se, but rather isolation (Massey & Denton, 1988). As such, it is able to

³ We did not discriminate between professional and non-professional secondary degrees

account for between unit variations in racial population proportions.

Missing county imputation. A major benefit of MrP approaches to sub-national estimation is that it affords an educated guess even for unobserved regions. However, to simplify our computational procedures, for each of the 156 missing counties, we used our county-level post-stratification tables to impute the most likely *type* of person for each county. Using a state-level hierarchical model estimated from the full data, we then predicted IAT D scores for this type of person, which was functionally a type of model-based missing data imputation.

Disaggregation. We obtained disaggregated estimates of D score $y_{u[l]}$ for county l located in region u simply by calculating the sample mean for observations from county l :

$$\bar{y}_{u[l]}^1 = \frac{\sum_i^{N_l} y_{u[l],i}}{N_l} \quad (18)$$

where $y_{u[l],i}$ represents response i from county $u[l]$, N_l represents the sample size for county $u[l]$, and $\bar{y}_{u[l]}^1$ is the estimated mean for county $u[l]$ obtained via method 1, disaggregation.

Post-stratification by age. To obtain post-stratified county D score estimates, we selected age as our post-stratification variable, as it is commonly used in applications of post-stratification and it has been used recently to derive county-level estimates of implicit racial bias from Project Implicit data (Leitner et al., 2016a,b; Hehman et al., 2017)⁴.

Specifically, we first estimated post-stratum sample means for each level of age nested in each county. Second, estimated county means as the weighted average of the post-stratum means, where weights corresponded to the county population proportion associated with each level of age. Given that post-stratification requires at least one observation for each post-stratum cell, we excluded counties that did not meet this

⁴ Note, though, that these studies collated data over a much larger time-span (2002-2013) and thus focused on data with different patterns of sparsity and non-representativeness. Thus, we do not mean to imply that our age post-stratified estimates are necessarily comparable to the estimates obtained in those studies.

requirement, which yielded a set of age post-stratified county D score estimates for 1,101 counties.

i.i.d. MrP - ML. County D score estimates obtained via an i.i.d. Maximum Likelihood MrP model⁵ were generated using a response model that included race, age, and gender as post-stratification variables, exposure (B) and county-level democratic vote share (Vote) as contextual factors, and county and region as geographic indicators:

$$\hat{y}_{u[l],i}^3 = \beta_0 + \beta_1 B + \beta_2 Vote + \alpha_{r[i]} + \alpha_{a[i]} + \alpha_{g[i]} + \alpha_u[l] \quad (19)$$

where random demographic effects for race, $\alpha_{r[i]}$, age, $\alpha_{a[i]}$, and gender $\alpha_{g[i]}$ were estimated as:

$$\begin{aligned} \alpha_r &\sim \mathcal{N}(0, \sigma_r^2), \text{ for } r = 1, \dots, R \\ \alpha_a &\sim \mathcal{N}(0, \sigma_a^2), \text{ for } a = 1, \dots, A \\ \alpha_g &\sim \mathcal{N}(0, \sigma_g^2), \text{ for } g = 1, \dots, G \end{aligned} \quad (20)$$

and random geographic effects were estimated for county, $\alpha_u[l]$, and region, α_u :

$$\begin{aligned} \alpha_u[l] &\sim \mathcal{N}(u, \sigma_l^2), \text{ for } l = 1, \dots, L \\ \alpha_u &\sim \mathcal{N}(0, \sigma_u^2), \text{ for } u = 1, \dots, U \end{aligned} \quad (21)$$

.

Next, using the trained model, predictions $\pi_u[l], j$ were made for each cross-classification j of race, age, and gender within each county $u[l]$.

Finally, the post-stratification step was implemented using the county-level population joint distribution for race, age, and gender estimated by the US Census:

⁵ This model was estimated using in the R Studio using the lme4 package (Bates et al., 2015)

$$\bar{Y}_{u[l]}^3 = \sum_j P_{u[l],j} \pi_{u[l],j} \quad (22)$$

i.i.d. MrP - INLA. To obtain i.i.d estimates using INLA, $\bar{Y}_{u[l]}^3$, we followed the same procedure as for model 3. However, estimation was conducted using the R-INLA package Version 18.07.12, (Lindgren & Rue, 2015) which offers an efficient and scalable implementation of Integrated Nested Laplace Approximation (Martino & Rue, 2009). For this model, weakly informative parameters (Gelman et al., 2013) were placed on all parameters ⁶:

$$\begin{aligned} \beta &\sim \mathcal{N}(0, 3) \\ \sigma_{[r,a,g]} &\sim \exp(1) \\ \sigma &\sim \exp(.1) \end{aligned} \quad (23)$$

MrP with spatial smoothing - INLA. To induce spatial smoothing, we extend Model 4 by altering the random effect of county, $\alpha_u[l]$, from an i.i.d. random effect to a spatially structured modified BYM effect, which accounts for both unstructured and structured variance (See section for discussion). Per Riebler et al. (2016), we place the following priors on $\alpha_u[l]$ and ϕ :

$$\begin{aligned} \alpha_u[l] &\sim \mathcal{N}(0, 5) \\ \phi &\sim \text{beta}(0.5, 0.5) \end{aligned} \quad (24)$$

i.i.d. MrsP - ML. MrsP estimates obtained from an i.i.d ML model were obtained by, first, extending the baseline county-level population joint distribution for gender, age,

⁶ Here, we report priors on standard deviates. However, R-INLA expects priors on precision, which is the inverse of variance ($\frac{1}{\text{sigma}^2}$); thus, we use simulation to convert the reported priors to distributions appropriate for R-INLA.

and race to also include education and ideology and then proceeding with model estimation, cross-classification prediction, and post-stratification.

To extend the baseline demographic joint distribution to include education, we followed the procedure for adjusted MrSP. This involved using data from the Religious Landscape Survey (Center, 2014) to estimate the joint distribution of race, age, gender, and education at the national level and then use the county-level marginal population distribution of education, obtained from the US Census, to *adjust* the estimated national joint distribution so that its *marginal* distribution of education matched a given county's marginal education distribution. Finally, we use the true county-level joint distribution of race, age, and gender and the relative weights obtained via the adjustment process in order to estimate the county-level joint distribution distribution of race, age, gender, and education.

Because county-level marginal distributions of ideology are not available, we followed an alternative procedure in order to add it to the 4-way joint distribution (Kastellec et al., 2015; Claassen & Trautmüller, 2018). Specifically, we constructed intermediate MrP models in order to generate estimates of the joint distribution of ideology, given the cross-classification of demographics (race, age, gender, and education), contextual factors (democratic vote share), and geographic indicators (county and region). That is, we applied MrP to the problem of estimating the mean probability of a person being conservative, moderate, or liberal, given the components included in this MrP model.

Further, in order to better explain demographic variance in ideology, we developed an additional intermediate MrP model for religious attendance (seldom or never attend, attend yearly to monthly, attend weekly or more), which we use as an indicator of religiosity. Thus, prior to estimating political ideology, we first extended the 4-way joint distribution to include frequency of religious attendance. Then, using a MrP model that included both education and religious attendance, as well as the baseline demographic variables, we

estimated the joint distribution of these variables to include political ideology.⁷

This multi-step procedure yielded a synthetic estimate of the joint distribution of race, age, gender, education, and political ideology. To obtain MrsP estimates, $\hat{y}_{u[l],i}^5$, over the joint distribution of these variables, they were included in a model analogous to Model 3. However, during estimation, model convergence was not achieved; the cause of this behavior was identified as the low variance of the random effects of education and this variable was subsequently dropped from the model. Thus, we obtained MrsP estimates using the following i.i.d. response model:

$$\hat{y}_{u[l],i}^5 = \beta_0 + \beta_1 B + \beta_2 Vote + \alpha_{r[i]} + \alpha_{a[i]} + \alpha_{g[i]} + \alpha_{p[i]} + \alpha_u[l] \quad (25)$$

where $\alpha_{p[i]}$ represents the random effect of political ideology. Using these estimates, the prediction and post-stratification steps were implemented as discussed above.

i.i.d. MrsP - INLA. This model was identical to model 4, aside from the addition of political ideology as a demographic predictor and post-stratification variable.

i.i.d. MrsP - modified BYM. This model was identical to model 7, aside from the modification of the i.i.d. random effect of county to a modified BYM effect.

Results

Across all Mr(s)P models, demographic, contextual, and geographic factors were estimated to have consistent but small effects on respondent's D scores, which is consistent with previous findings (Greenwald et al., 2003). Specifically, both exposure and democratic vote share were associated with small D score decreases (See Table 5) and while D scores

⁷ In order to reduce computational cost, the probabilities for each category of religious attendance and political ideology were estimated using a series of logistic regression models. For example, to estimate religious attendance, we first estimated a model discriminating between seldom-never and yearly-monthly respondents. Then, we estimated a separate model discriminating between the union of these respondents and weekly+ respondents.

varied across demographic groups, this variation ranged from approximately 0.02 to 0.07 standard deviations.

	i.i.d MrP	i.i.d MrsP
Intercept	0.32*** (0.03)	0.33*** (0.05)
Vote	-0.02*** (0.00)	-0.01*** (0.00)
B	-0.04*** (0.00)	-0.05*** (0.00)
Random effects		
County	0.032	0.029
Division	0.019	0.02
Race	0.017	0.015
Gender	0.038	0.029
Ideology		0.068
Residual	0.435	0.432

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 5

MrP ML response models

Further, county-level estimates obtained from the Mr(s)P methods were strongly correlated ($rs > 0.98$, $\tau > 0.92$) and generally consistent, such that differences between estimates fell between ± 0.05 . However, we did also observe that i.i.d. MrP estimates, on average, were approximately 0.02 points lower than MrsP and spatially smoothed MrP estimates, which indicates that estimates obtained via the latter methods were slightly more robust to shrinkage than those obtained via MrP. Further, comparison of the disagreement between the i.i.d. and spatially smoothed estimates among counties with

small samples and their disagreement for all counties indicated that spatial smoothing slightly impacted the rank ordering of county means for counties with small samples. Specifically, while MrP and spatially smoothed MrP estimates for counties with fewer than 5 samples had a Kendall’s τ rank order correlation (Kendall, 1938) of $\tau = 0.89$, their estimates for all counties had a rank order correlation of $\tau = 0.94$. Notably, no such effect of sample size on τ was observed in comparisons of i.i.d. MrP and MrsP estimates.

Nonetheless, while estimates obtained from the targeted MrP methods deviated in expected ways, these deviations were minor. Accordingly, for the remainder of these analyses, we focus on the disagreement between disaggregation and post-stratification estimates and estimates obtained via ML MrP. As all findings for this comparison are consistent across comparisons between disaggregation and post-stratification estimates and estimates obtained from the other MrP approaches, this allows us to simplify the presentation of our results.

As expected, estimates derived via disaggregation and post-stratification by age deviated strongly from estimates obtained via our implementations of MrP. Among counties for which age post-stratified estimates could be calculated, estimates correlated only moderately with MrP estimates we observed Pearson and Kendall’s correlations of $r = 0.50$, $\tau = 0.39$. Further, these correlations only increased modestly when counties with low sample sizes were excluded. For example, even among counties with at least 50 samples, MrP and age post-stratified estimates were only correlated at $r = 0.55$ and $\tau = 0.43$.

Regarding estimates obtained via disaggregation, estimated across the entire set of counties, we observed correlations between disaggregation and MrP estimates of $r = 0.49$ and $\tau = 0.37$. However, among counties with more than 25 observations (approximately 46% of the total N of counties), this correlation increased to $r = 0.81$ and $\tau = 0.64$. Further, among counties with more than 100 observations (approximately 23% of the total N of counties), this correlation increased further to $r = 0.94$ and $\tau = 0.81$.

However, examination of the distribution of disagreement shows that while

disagreement shrinks as the county sample size cut-off increases, this reduction in disagreement requires a dramatic decrease in the number of counties eligible for analysis (See Table 6). These results are also depicted in Figure 3, which shows the density of disagreement scores for counties within various sample-size bins as well as the linear relationship between the estimates obtained from disaggregation and MrP.

Table 6

Summary of disagreement between MrP and disaggregation

Minimum N (% of counties)	Mean	Median	SD	Min.	1st Qu.	3rd Qu.	Max.
0 (100%)	0	0	0.16	-1.40	-0.05	0.070	0.92
5 (78%)	-0.01	-0.01	0.10	-0.43	-0.04	0.03	0.600
25 (47%)	-0.01	-0.01	0.05	-0.23	-0.03	0.01	0.22
50 (34%)	-0.01	-0.01	0.03	-0.18	-0.02	0	0.22
100 (23%)	-0.01	-0.01	0.02	-0.18	-0.02	0	0.090

Summary of disagreement between MrP and disaggregation estimates over different minimum cut-offs for sample size.

In order to visualize potential differences in spatial patterning, we also generated county-level D score maps using estimates obtained via disaggregation and MrP (See Figure 4). Comparison of these maps highlight several key observations. For example, whereas the disaggregation estimates range from approximately -1 to 1.3, MrP estimates range from approximately 0 to 0.5. Further, the MrP estimates appear to show a more reasonable spatial structuring, such that mean levels of bias shift gradually across space. In contrast, the disaggregated estimates show rather wild fluctuation, even within small areas.

While the D score maps also facilitate approximate evaluations of *where* disaggregation and MrP disagree most strongly, to better visualize their disagreement, we

also generated a map depicting the geographic distribution of their absolute difference scores. Further, to better understand this distribution, we also constructed a map of county sample sizes 5. Together, these facilitate approximate evaluation of where and to what extent disagreement is driven by sample size. Notably, while there is a clear association between disagreement and sample size, this association is not particularly strict. Indeed, the correlation between the absolute difference between MrP and disaggregation estimates and county sample size is only -0.18 .

Figure 3. Comparison of disaggregated and MrP estimates of county-level implicit racial bias. In both panels, color indicates a range of observed county sample sizes (the number of counties that fall within each range is indicated parenthetically in the legend). The bottom panel plots county predictions obtained from MrP (x-axis) against those obtained via disaggregation (y-axis). The fit-line indicates perfect agreement. The top panel depicts the density (y-axes) of the differences (x-axes) between disaggregation and MrP predictions as a function of observed sample size. Note that the y- and x-axis scales in the top panel are not constant.

Figure 4. County level race IAT D score estimates obtained via disaggregation and MrP

Discussion

In this study, we pursued a pragmatic comparison of estimates obtained via by applying disaggregation and a range of approaches to Mr(s)P to a large, opt-in online study. Specifically, we focused on estimating county-level implicit racial bias against Blacks, measured via the IAT and represented as a D score. Our results highlight a number of important issues related to estimating sub-national scores of psychological constructs from large but non-random data.

First, it is notable that across the Mr(s)P models, all the fixed and random effects had relatively small magnitudes and considerable residual variance was left unexplained. As discussed early, the power of MrP depends considerably on the magnitude of its effects and the capacity of the response model to explain variance in the outcome. The fact that we were not able to develop strong response models (e.g. with large demographic or contextual effects) potentially highlights the difficulty of working with psychological constructs, which may yield noisier measurements compared to the kinds of political preference outcomes typically estimated using MrP approaches.

However, despite the relative weakness of our response models, it is essential to note that all of our MrP estimates deviated dramatically from estimates obtained via both

Figure 5. County sample size and disagreement between disaggregation and MrP estimates

disaggregation and post-stratification. While post-stratification should, in theory, improve estimation compared to disaggregation, we suspect that here it was inhibited due to the combination of post-stratum cell sparsity and the inability to adjust for variation as a function of other demographic variables (Little, 1993; Park et al., 2004).

Further, while county disagreement between disaggregation and MrP was related to county sample size, the correlation was relatively weak ($r = -0.18$). Consistent with Study 1, this suggests that, at least with this data, increasing sample size does not sufficiently address potential sources of bias. This is an important caveat to previous findings that MrP and disaggregation quickly converge as sample size increases. For this data, and we suspect most data generated via similar sampling mechanisms, increasing sample size *does* induce convergence, but this rate of convergence is dramatically slower than observed in studies conducted with representative data.

Overall, given our previous studies and the literature on MrP, we view these MrP estimates as the best possible guess, given our response models, regarding county-level racial bias. In addition to addressing potential problems of representativeness, these MrP approaches facilitate informed and conservative estimation for counties with low sample sizes. This can have a major impact on downstream analyses, as it enables researchers to

avoid excluding entire geographic regions due to insufficient responses. While such exclusion practices may not necessarily have deleterious effects on inference, in situations where a downstream outcome (e.g. outcomes hypothesized to be associated with county-level racial bias) is correlated with non-response, excluding regions with low response rates will lead to biased inferences. While MrP cannot, of course, offer guaranteed estimates for areas with few observations — or areas with any number of observations, for that matter — the capacity to generate informed predictions for these values enables researchers to conduct sensitivity analyses by estimating models that both include and exclude low N regions.

Finally, our Study 3 results highlight the fact that a more complex MrP model will not necessarily yield substantively different estimates. In our case, the combination of low residual spatial auto-correlation and small demographic effects minimized the difference between MrP estimates and estimates obtained with spatial smoothing or synthetic joint distributions. Under different conditions, these models likely would have yielded more divergent estimates. Nonetheless, these results highlight the fact that, at least under some conditions, the simplest MrP procedure can yield estimates that are comparable to procedures that are far more computationally complex. This is important, as it suggests that simpler approaches can reasonably be employed for model development and exploratory analyses.

General Discussion

One of the primary difficulties for psychological research is establishing connections between hypothesized constructs and real-world phenomena. While psychologists are experts at simulating phenomena in laboratory settings and developing indirect (e.g. survey based) measures of target phenomena, establishing external validity remains one of the primary challenges for psychological research. While geographic studies of psychological phenomena raise their own challenges, they also directly supplement

conventional approaches to psychological research. More specifically, geographic approaches to psychological research offer an opportunity to directly investigate the association between psychological constructs and real-world outcomes (Rentfrow & Jokela, 2016).

This is, of course, not a revelation, as cultural psychology pioneered cross-cultural (e.g. international) studies decades ago. However, early investigations of the geographic study of psychological outcomes were limited by focus on nation-level variation. Recently, however, researchers have begun focusing psychological constructs' sub-national variation and associations with target outcomes. This work has yielded a number of important findings, such as new evidence for the deleterious effects of racial prejudice (Leitner et al., 2016a; Hehman et al., 2017; Orchard & Price, 2017; Rae et al., 2015) and associations between personality and various target outcomes (Jokela et al., 2015; Rentfrow et al., 2015; McCann, 2017b,a, 2018; Götz et al., 2018), as well as insight into the spatially structured patterning of personality (Rentfrow et al., 2015, 2008, 2013; Jokela et al., 2015) and moral values (Hoover et al., 2018).

However, while the sub-national study of psychological constructs affords a range of exciting new opportunities for research, it also raises a number of challenges that are not commonly encountered or addressed in psychological research. In this work, our goal was to offer a comprehensive introduction to state-of-the-art methods, Mr(s)P, for addressing these challenges that have been developed in other fields. Beyond this, we sought to evaluate the performance of these methods under conditions similar to those in which many geographic studies of psychological phenomena have been conducted. Specifically, the methods reviewed in this work were both designed for and have been most often applied to relatively small, randomized, and representative samples. In contrast, psychologists today often find themselves working with large, non-random, and non-representative samples with non-uniform geo-graphic sparsity. Accordingly, in order to provide researchers with an informed introduction to these methods, we addressed several questions regarding the optimal approach to sub-national estimation under such conditions.

In Study 1, we evaluated the differential performance of disaggregation — a simple but widely used approach to small-area estimation — and MrP under varying conditions of sample size and response bias. This study provided strong evidence that under conditions of response bias, MrP dramatically outperforms disaggregation, regardless of sample size. In Study 2, we extended recent research on an a modified approach to MrP estimation (Selb & Munzert, 2011), which relies on local spatial smoothing to improve estimates. While previous work found little benefit of local spatial smoothing (Selb & Munzert, 2011; Hanretty et al., 2016), our simulations suggest that spatial smoothing can substantively improve estimates for regions with small sample sizes. Given that much of geographic psychological research faces issues with geographic sparsity, these findings suggest a relevant and actionable approach to improving sub-national estimates.

Finally, in Study 3, we explore divergences among county-level estimates of implicit racial bias derived by applying a variety of estimation procedures to a large dataset generated via an opt-in online sampling mechanism. This study raised several important points regarding sub-national estimation with large, non-random and non-representative psychological data. Most importantly, we found that while all sets of estimates obtained via a MrP method showed strong agreement, disaggregation and post-stratification by age yielded substantively different estimates. While we cannot evaluate the accuracy of any one set of estimates, given the literature on MrP as well as our findings in Studies 1 and 2, we interpret these results as suggesting that psychological researchers should strongly consider including some version of MrP in their sub-national estimation procedures.

However, Study 3's results also indicate that, under some conditions, simple implementations of MrP may yield estimates that are largely equivalent to estimates generated from more sophisticated implementations. While MrP implementations should be chosen based on context, a priori knowledge about the data generating process, and pragmatic limitations, these results suggest that simple approaches to MrP can offer a desirable balance between performance and computational cost.

Overall, given the relative ease of implementing MrP and the improvements in estimation accuracy and stability that it offers, we see little reason for this method to not be more widely applied to sub-national geographic studies of psychological phenomena. While today's large, online, opt-in samples have opened many new opportunities for studying the geographic distribution of psychological constructs, it is important that limitations of these samples are addressed as well as is possible. In our view, MrP useful tool that can help psychological researchers work toward this goal.

References

- Allik, J., Realo, A., Mõttus, R., Pullmann, H., Trifonova, A., McCrae, R. R., & and 56 Members of the Russian Character and Personality Survey. (2009, November). Personality traits of russians from the observer's perspective. *European journal of personality*, 23(7), 567–588.
- Banerjee, S., Carlin, B. P., & Gelfand, A. E. (2014). *Hierarchical modeling and analysis for spatial data, second edition*. CRC Press.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. doi: 10.18637/jss.v067.i01
- Beale, C. M., Lennon, J. J., Yearsley, J. M., Brewer, M. J., & Elston, D. A. (2010, February). Regression analysis of spatial data. *Ecology letters*, 13(2), 246–264.
- Besag, J., & Kooperberg, C. (1995, December). On conditional and intrinsic autoregressions. *Biometrika*, 82(4), 733–746.
- Besag, J., York, J., & Mollié, A. (1991, March). Bayesian image restoration, with two applications in spatial statistics. *Annals of the Institute of Statistical Mathematics*, 43(1), 1–20.
- Blangiardo, M., & Cameletti, M. (2015a). *Spatial and spatio-temporal bayesian models with R - INLA*. Wiley.
- Blangiardo, M., & Cameletti, M. (2015b). *Spatial and spatio-temporal bayesian models with R - INLA*.
- Bureau, U. C. (n.d.). *2015 Census Bureau's MAF/TIGER geographic database*.
- Buttice, M. K., & Highton, B. (2013). How does multilevel regression and poststratification perform with conventional national surveys? *Political analysis: an annual publication of the Methodology Section of the American Political Science Association*.

Center, P. R. (n.d.). *2011-2015 american community survey 5-year estimates*.

Center, P. R. (2014). *2014 religious landscape study*. Washington, D.C..

Claassen, C., & Traunmüller, R. (2018, April). Improving and validating survey estimates of religious demography using bayesian multilevel models and poststratification. *Sociological methods & research*, 0049124118769086.

Deville, J.-C., Särndal, C.-E., & Sautory, O. (1993, September). Generalized raking procedures in survey sampling. *Journal of the American Statistical Association*, 88(423), 1013–1020.

Eberly, L. E., & Carlin, B. P. (2000). Identifiability and convergence issues for markov chain monte carlo fitting of spatial models. *Statistics in medicine*, 19(17-18), 2279–2294.

Erikson, R. S., Wright, G. C., & McIver, J. P. (1993). *Statehouse democracy: Public opinion and policy in the american states*. Cambridge University Press.

Fowler, L. (2016, March). The states of public opinion on the environment. *Environmental politics*, 25(2), 315–337.

Gelman, A. (2007, May). Struggles with survey weighting and regression modeling. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 22(2), 153–164.

Gelman, A. (2014, February). How bayesian analysis cracked the Red-State, Blue-State problem. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 29(1), 26–35.

Gelman, A. (2018, August). The most important aspect of a statistical analysis is not what you do with the data, it's what data you use (survey adjustment edition) [Blog post]. <http://andrewgelman.com/2018/08/07/important-aspect-statistical-analysis-not-data-data-use-survey-adjustment-edition/>. (Accessed: 2018-8-8)

- Gelman, A., & Hill, J. (2006). *Data analysis using regression and Multilevel/Hierarchical models*. Cambridge University Press.
- Gelman, A., & Little, T. C. (1997). Poststratification into many categories using hierarchical logistic regression. *Survey methodology*, 23(2), 127–135.
- Gelman, A., Stern, H. S., Carlin, J. B., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis*. Chapman and Hall/CRC.
- Götz, F. M., Ebert, T., & Rentfrow, P. J. (2018). Regional cultures and the psychological geography of switzerland: Person–Environment–Fit in personality predicts subjective wellbeing. *Frontiers in psychology*, 9, 517.
- Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003, August). Understanding and using the implicit association test: I. an improved scoring algorithm. *Journal of personality and social psychology*, 85(2), 197–216.
- Greenwald, A. G., Poehlman, T. A., Uhlmann, E. L., & Banaji, M. R. (2009, July). Understanding and using the implicit association test: III. meta-analysis of predictive validity. *Journal of personality and social psychology*, 97(1), 17–41.
- Guélat, J., & Kéry, M. (2018, May). Effects of spatial autocorrelation and imperfect detection on species distribution models. *Methods in ecology and evolution / British Ecological Society*, 4, 581.
- Hackett, C., Grim, B., Stonawski, M., Skirbekk, V., Kuriakose, N., Potančoková, M., . . . Zurlo, G. A. (2014). Methodology of the pew research global religious landscape study. *Yearbook of International Religious Demography 2014*, 167.
- Hanretty, C., Lauderdale, B. E., & Vivyan, N. (2016). Comparing strategies for estimating constituency opinion from national survey samples. *Political Science Research and Methods*, 1–21.

- Hehman, E., Flake, J. K., & Calanchini, J. (2017, July). Disproportionate use of lethal force in policing is associated with regional racial biases of residents. *Social psychological and personality science*, 1948550617711229.
- Holt, D., & Smith, T. M. F. (1979). Post stratification. *Journal of the Royal Statistical Society. Series A*, 142(1), 33–46.
- Hoover, J., Zhao, C., & Dehghani, M. (2018). *MapYourMorals :: USC CSSL*. <https://mapyourmorals.usc.edu/#/>. (Accessed: 2018-10-10)
- Howe, P. D., Mildenerger, M., Marlon, J. R., & Leiserowitz, A. (2015, April). Geographic variation in opinions on climate change at state and local scales in the USA. *Nature climate change*, 5(6), 596–603.
- Jokela, M., Bleidorn, W., Lamb, M. E., Gosling, S. D., & Rentfrow, P. J. (2015, January). Geographically varying associations between personality and life satisfaction in the london metropolitan area. *Proceedings of the National Academy of Sciences of the United States of America*, 112(3), 725–730.
- Kalton, G., & Flores-Cervantes, I. (2003). Weighting methods. *Journal of official statistics*, 19(2), 81.
- Kastellec, J. P., Lax, J. R., Malecki, M., & Phillips, J. H. (2015). Polarizing the electoral connection: Partisan representation in supreme court confirmation politics. *The journal of politics*, 77(3), 787–804.
- Kendall, M. G. (1938). A new measure of rank correlation. *Biometrika*, 30(1/2), 81–93.
- Krimmel, K. L., Lax, J. R., & Phillips, J. H. (2012). *Public opinion and gay rights: Do members of congress follow their constituents' preferences?* (Tech. Rep.). Working Paper.

- Lax, J. R., & Phillips, J. H. (2009, January). How should we estimate public opinion in the states? *American journal of political science*, 53(1), 107–121.
- Lee, D. (2011, June). A comparison of conditional autoregressive models used in bayesian disease mapping. *Spatial and spatio-temporal epidemiology*, 2(2), 79–89.
- Leemann, L., & Wasserfallen, F. (2017, October). Extending the use and prediction precision of subnational public opinion estimation: EXTENDING USE AND PRECISION OF MrP. *American journal of political science*, 61(4), 1003–1022.
- Leitner, J. B., Hehman, E., Ayduk, O., & Mendoza-Denton, R. (2016a). Blacks’ death rate due to circulatory diseases is positively related to whites’ explicit racial bias: A nationwide investigation using project implicit. *Psychological science*, 27(10), 1299–1311.
- Leitner, J. B., Hehman, E., Ayduk, O., & Mendoza-Denton, R. (2016b, October). Racial bias is associated with ingroup death rate for blacks and whites: Insights from project implicit. *Social science & medicine*.
- Lindgren, F., & Rue, H. (2015). Bayesian spatial modelling with R-INLA. *Journal of statistical software*, 63(19).
- Little, R. J. A. (1993). Post-Stratification: A modeler’s perspective. *Journal of the American Statistical Association*, 88(423), 1001.
- Martino, S., & Rue, H. (2009). Implementing approximate bayesian inference using integrated nested laplace approximation: A manual for the inla program. *Department of Mathematical Sciences, NTNU, Norway*.
- Massey, D. S., & Denton, N. A. (1988, December). The dimensions of residential segregation. *Social forces; a scientific medium of social study and interpretation*, 67(2), 281–315.

- McCann, S. J. H. (2017a, December). Higher USA state resident neuroticism is associated with lower state volunteering rates. *Personality & social psychology bulletin*, 43(12), 1659–1674.
- McCann, S. J. H. (2017b, December). The relation of state resident neuroticism levels to state cancer incidence in the USA. *Current psychology*.
- McCann, S. J. H. (2018, April). U.S. state resident big five personality and work satisfaction: The importance of neuroticism. *Cross-cultural research: official journal of the Society for Cross-Cultural Research / sponsored by the Human Relations Area Files, Inc*, 52(2), 155–191.
- Mooney, C. Z. (1997). *Monte carlo simulation*. SAGE Publications.
- Orchard, J., & Price, J. (2017, May). County-level racial prejudice and the black-white gap in infant health outcomes. *Social science & medicine*, 181, 191–198.
- Park, D. K., Gelman, A., & Bafumi, J. (2004). Bayesian multilevel estimation with poststratification: State-Level estimates from national polls. *Political analysis: an annual publication of the Methodology Section of the American Political Science Association*, 12(4), 375–385.
- Pew. (2018). Demographics of Internet Users. Pew Internet & American Life Project. (<http://www.pewinternet.org/fact-sheet/internet-broadband/>)
- Pew Research Center. (2018, January). For Weighting Online Opt-In Samples, What Matters Most? <http://www.pewresearch.org/2018/01/26/for-weighting-online-opt-in-samples-what-matters-most/>. (Accessed: 2018-8-8)
- Rae, J. R., Newheiser, A.-K., & Olson, K. R. (2015, July). Exposure to racial Out-Groups and implicit race bias in the united states. *Social psychological and personality science*, 6(5), 535–543.

- Raudenbush, S. W., & Bryk, A. S. (2002). *Hierarchical linear models: Applications and data analysis methods* (Vol. 1). Sage.
- Rentfrow, P. J., Gosling, S. D., Jokela, M., Stillwell, D. J., Kosinski, M., & Potter, J. (2013, December). Divided we stand: three psychological regions of the united states and their political, economic, social, and health correlates. *Journal of personality and social psychology*, 105(6), 996–1012.
- Rentfrow, P. J., Gosling, S. D., & Potter, J. (2008, September). A theory of the emergence, persistence, and expression of geographic variation in psychological characteristics. *Perspectives on psychological science: a journal of the Association for Psychological Science*, 3(5), 339–369.
- Rentfrow, P. J., & Jokela, M. (2016, December). Geographical psychology: The spatial organization of psychological phenomena. *Current directions in psychological science*, 25(6), 393–398.
- Rentfrow, P. J., Jokela, M., & Lamb, M. E. (2015, March). Regional personality differences in great britain. *PloS one*, 10(3), e0122245.
- Riebler, A., Sørbye, S. H., Simpson, D., & Rue, H. (2016, August). An intuitive bayesian spatial model for disease mapping that accounts for scaling. *Statistical methods in medical research*, 25(4), 1145–1165.
- Selb, P., & Munzert, S. (2011). Estimating constituency preferences from sparse survey data using auxiliary geographic information. *Political analysis: an annual publication of the Methodology Section of the American Political Science Association*.
- Simpson, D., Rue, H., Riebler, A., Martins, T. G., & Sørbye, S. H. (2017, February). Penalising model component complexity: A principled, practical approach to constructing priors. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 32(1), 1–28.

- Steenbergen, M. R., & Jones, B. S. (2002). Modeling multilevel data structures. *American journal of political science*, 46(1), 218–237.
- Wang, W., Rothschild, D., Goel, S., & Gelman, A. (2015). Forecasting elections with non-representative polls. *International journal of forecasting*, 31(3), 980–991.
- Xu, F. K., Lofaro, N., Nosek, B. A., & Greenwald, A. G. (2013, September). *Race IAT 2002-2017*. <https://osf.io/52qxl/>. OSF. (Accessed: 2018-9-30)