

UC Riverside

UC Riverside Previously Published Works

Title

Regional intergroup bias

Permalink

<https://escholarship.org/uc/item/661509jq>

Authors

Calanchini, Jimmy

Helman, Eric

Ebert, Tobias

et al.

Publication Date

2022

DOI

10.1016/bs.aesp.2022.04.003

Peer reviewed

THIS MANUSCRIPT HAS BEEN ACCEPTED FOR PUBLICATION IN
Advances in Experimental Social Psychology, vol. 66.

Regional Intergroup Bias

Jimmy Calanchini¹, Eric Hehman², Tobias Ebert³, Emily Esposito^{1*}, Deja Simon^{1*}, & Liz Wilson^{1*}

¹Department of Psychology, University of California, Riverside, Riverside, CA 92521

²Department of Psychology, McGill University, Montreal, QC, H3A 0G4

³Mannheim Centre for European Social Research, University of Mannheim 68159

*These authors contributed equally to the preparation of this manuscript, and are listed alphabetically.

Correspondence to: Jimmy Calanchini, Department of Psychology, University of California, Riverside, E-mail: jimmy.calanchini@ucr.edu

Author Contributions: J.C., E.H. and T.E. developed the ideas and wrote the manuscript. E.E, D.S., and L.W curated the data and analyses; E. H., T.E., and L.W. created the figures.

Funding Disclosure: This work was partially supported by a postdoctoral fellowship from the Alexander von Humboldt Foundation to J.C.

Acknowledgments: We thank Tony Greenwald and Kate Ratliff for their helpful feedback on this manuscript.

Word Count (including figures, captions, and references): 21,795

Abstract

Recent advances in large-scale data collection have created new opportunities for psychological scientists who study intergroup bias. By leveraging big data, researchers can aggregate individual measures of intergroup bias into regional estimates to predict outcomes of consequence. This small-but-growing area of study has already impacted the field with well-powered research identifying relationships between regional intergroup biases and societally-important, ecologically-valid outcomes. In this chapter, we summarize existing regional intergroup bias research and review relevant theoretical perspectives. Next, we present new and recent evidence that cannot be explained by existing theory, and offer a new perspective on regional intergroup bias that highlights aggregation as changing its' qualitative nature relative to individual intergroup bias. We conclude with a discussion of some of the important challenges that regional intergroup bias research will need to address in moving forward, focusing on issues of prediction and causality; constructs, measures, and data sources; and levels of analysis.

Abstract word count: 150

Keywords: intergroup bias, attitudes, stereotypes, prejudice, regional psychology, macropsychology

Recent advances in large-scale data collection have created new opportunities for psychological scientists who study intergroup bias. Some disciplines, such as sociology, economics, and political science, have tested hypotheses related to intergroup bias with region-level data for many decades. However, until relatively recently, psychological data on intergroup bias were primarily collected through small, controlled experiments in laboratories on university campuses. Technology now facilitates the collection of massive amounts of data from diverse populations and locations. Whereas previously, collecting data across communities, states, or nations often required significant time, logistics, and resources (e.g., to send teams of research assistants into the field), data can now be collected online at large scales for the relatively low costs of developing, maintaining, and promoting a website. The internet has enabled psychology-related initiatives such as Project Implicit, YourMorals, and the MyPersonality Project to collect data from millions of interested visitors. Intergroup bias researchers can leverage newly-available datasets like these to aggregate the responses of individuals into regional estimates and, in conjunction with other sources of data, predict outcomes of consequence.

This small-but-growing area of study has already had an impact by identifying relationships between regional intergroup biases and important, real-world behavioral outcomes. For example, relative to local demographics, Black but not White people are disproportionately killed by police in areas that are more racially prejudiced (Hehman et al., 2018). Black but not White babies who are born in racially prejudiced regions are more likely to be born prematurely and have low birth weight (Orchard & Price, 2017). This regional approach to studying intergroup bias offers several advantages over traditional laboratory-based approaches. By aggregating the responses of geographically-proximate individuals on psychological measures of bias – which can include traditional survey instruments or indirect measures of bias such as the

implicit association test (IAT: Greenwald et al., 1998) – the regional approach allows systematic study of outcomes related to intergroup bias that are important but happen relatively infrequently or are otherwise difficult to isolate in the laboratory. Additionally, regional intergroup bias research aligns well with new norms in scientific best practices. Regional analyses often rely upon datasets that reflect very large and diverse samples, which afford greater statistical validity and generalizability than do most laboratory-based experiments. These datasets are often publicly available, which facilitates transparency and open science. Moreover, whereas laboratory-based intergroup bias research often contrives situations in order to study judgments and behaviors, regional intergroup bias research instead examines the real-world outcomes that laboratory-based research aspires to recreate. At the same time, regional intergroup bias research has demonstrated that psychological constructs provide explanatory power beyond the sociodemographic factors classically examined by other disciplines adopting a regional approach (e.g., race, socioeconomic status, population density). In this capacity, regional intergroup bias research brings a decidedly psychological perspective to bear on topics that have traditionally been the purview of other social sciences.

Studies involving regional intergroup bias have now proliferated such that a summary and synthesis will further advance this area of study. Consequently, we begin this chapter with a review of the existing literature on regional intergroup bias, followed by a discussion of theoretical perspectives that are related to regional intergroup bias. Next, we present new and recent evidence that cannot be explained by existing theory, offer a new perspective on regional intergroup bias, and conclude with discussion of some of the important challenges that this area of research will need to address in moving forward.

1. Empirical Evidence for Regional Intergroup Bias

The regional approach has been applied to predict a wide variety of outcomes related to intergroup bias, and the impact of this work can already be seen: At the time of writing, the 32 articles we summarize below have been cited over 2,585 times in total according to Google Scholar. Most of this research has been conducted in the context of Black/White racial bias in the United States, but some research has focused on biases in other intergroup domains and other geographic regions. We summarize here some examples of different contexts and domains in which regional intergroup bias has predicted outcomes.

1.1 Regional Racial Bias

Much of the regional research on Black/White racial bias highlights inequalities in health outcomes in the United States. Black people are less likely to have access to affordable health care and are more likely to die of circulatory disease in US counties in which White people are more racially biased (Leitner et al., 2016a; Zestcott et al., 2021). More generally, county-level racial biases of Black people are associated with Black mortality rates, and county-level racial biases of White people are associated with White mortality rates (Leitner et al., 2016b). Further, Black but not White mothers who live in counties with stronger racial biases suffer from more complicated pregnancies and more often give birth to low birth weight babies (Orchard & Price, 2017). Relative to the local population, Black but not White people are disproportionately killed by police in areas where White residents have higher levels of racial bias (Herman et al., 2018). Similarly, Black but not White drivers are disproportionately stopped by police in counties where White residents have higher levels of racial bias (Ekstrom et al., 2021; Stelter et al., 2021). Psychotherapy interventions are less effective for Black but not White youth who live in states that are more racially biased (Price et al., 2021a), and Black but not White youth also have smaller hippocampal volume in regions that are more racially biased (Hatzenbuehler et al.,

2021). HIV prevention interventions are less effective at improving condom use among Black people in regions where White people are more racially biased (Reid et al., 2014). Racial minorities disproportionately benefit from Medicaid programs, and states with higher levels of racial bias spend less per Medicaid enrollee (Leitner et al., 2018). Taken together, regional racial biases predict a wide variety of negative health outcomes for Americans – and, especially, Black Americans.

In addition to health disparities, regional racial bias has been linked to a wide variety of other outcomes and population characteristics in the United States. Black children are adopted from foster care at lower rates in states with higher levels of racial bias (Bell et al., 2021). In the context of education, Black students are disciplined more and perform worse on tests than White students in counties where teachers are more racially biased (Chin et al., 2020) and where residents are more racially biased (Riddle & Sinclair, 2019), and achievement gaps between Black and White students are larger in counties where residents are more racially biased (Pearman, 2021). Racial bias is also higher on college campuses with less diverse faculty (Vuletich & Payne, 2019), in counties where Black people have less economic mobility (Payne et al., 2019), and in states with greater prevalence of infectious disease (O’Shea et al., 2019), higher crime rates (Johnson & Chopik, 2019), higher levels of violence and economic inequality (Kunst et al., 2017), higher internet search rates for racial slurs (Hehman et al., 2019; Rae et al., 2015), higher proportions of Black people in the population (Rae et al., 2015), and where higher proportions of White residents endorse the sentiment that race receives “too much attention” in America today (Hehman et al., 2019).

1.2 Other Regional Intergroup Biases

Though most regional intergroup bias research to date has focused on Black/White racial bias specifically in the United States, some research has focused on biases in other intergroup domains and other geographic regions. Gender and sexuality biases, in particular, are linked to a variety of regional outcomes. Psychotherapy interventions are less effective for girls but not boys who live in more sexist communities (Price et al., 2021b), and middle school-aged girls lag behind their male peers in science achievement in countries with stronger career-related gender stereotypes (Nosek et al., 2009). Gay men and lesbian women in the United States are more likely to move away from regions where community members more strongly endorse anti-gay policies (Esposito & Calanchini, 2021). People with HIV in the New England area of the United States report higher levels of psychological distress and lower levels of physical well-being in communities with higher levels of bias against people with HIV (Miller et al., 2016). Much of this research focuses on disparities that negatively impact members of stigmatized groups, but recent work from our labs highlights one encouraging trend: Sexuality biases became more positive towards gay men and lesbians after legalization of same-sex marriage in the United States (Ofosu et al., 2019).

In addition to biases related to race, gender, and sexuality, a handful of other intergroup biases have been studied from a regional perspective. More people attended unity rallies after the Charlie Hebdo terrorist attacks in regions of France with lower average levels of anti-Arab bias (Zerhouni et al., 2016). Bias against Native Americans increased following the removal of Native American sports mascots, and especially in the American states where the mascots' teams resided (Jimenez et al., 2021). Asians are perceived to be less American in metropolitan areas where relatively fewer Asian Americans reside (Devos & Sadler, 2019; Devos et al., 2021).

Countries with higher proportions of overweight residents have higher average levels of bias in favor of thin versus overweight people (Marini et al., 2013).

Taken together, existing research on regional intergroup bias has examined a wide variety of outcomes, but at the same time this work has largely been clustered in one domain and in one context: Black/White racism in the United States. On the one hand, such a restricted range hinders the generalizability of findings from this body of research – a point we elaborate on later in this chapter. However, on the other hand, this relative focus can also be healthy for the field because it provides a basis for deeper investigations into the construct validity and psychometrics of intergroup bias as a regional construct.

1.3 The Construct Validity and Psychometrics of Regional Racial Bias

Psychology has traditionally conceptualized intergroup biases as existing within people (e.g., Allport, 1954). Consequently, measures of intergroup bias were developed and validated to be used at the individual level based on theory articulated at the individual level. However, and importantly, validity evidence for a measure of a construct is necessarily limited to a specific use or purpose (Kane, 2013), so new evidence is needed when an established measure is used in a new context, in order to assess whether the previously established interpretation is valid in the new context.

With this perspective in mind, we embarked on an investigation into the construct validity and psychometrics of regional racial bias (Hehman et al., 2019). As indices of construct validity, we focused on the predictive and discriminant validity of regional measures of racial bias. Predictive validity refers to the extent that an operationalization of a construct (in this case, regional racial bias) corresponds with theoretically-relevant outcomes, whereas discriminant validity refers to the extent that the operationalization does not correspond with theoretically-

unrelated outcomes. In a demonstration of predictive validity, we found that states with higher levels of racial bias also had higher proportions of internet searches that included racial slurs (see also Rae et al., 2015). The evidence reviewed above, in section 1.1 Regional Racial Bias, also demonstrates the predictive validity of regional racial bias. In a demonstration of discriminant validity, we found that state-level racial bias was unrelated to state-level birth rates.¹ Taken together, these examples demonstrate the construct validity of regional racial bias, in that it correlates with outcomes that are theoretically related to bias, but does not correlate with outcomes that are theoretically unrelated to bias.

To examine the psychometrics of regional racial bias, we investigated correspondence between different measures of regional racial bias. Racial bias – and intergroup bias more generally – can be measured in two main ways: explicitly and implicitly (Figure 1). Explicit measures rely on self-report through methods such as feeling thermometers, semantic differentials, or Likert-type response scales. In contrast, implicit measures such as the IAT (Greenwald et al., 1998) and affect misattribution procedure (AMP: Payne et al., 2005) infer biases through the speed or accuracy of responses, rather than from the contents of responses, *per se*.² Previous research has demonstrated that, at the individual level, measures of explicit bias generally demonstrate good retest reliability but measures of implicit bias generally demonstrate

¹ Other researchers have used a similar approach to demonstrate the discriminant validity of regional racial bias. Payne and colleagues (2019) found that regional variation in the proportion of enslaved population before the American Civil War reliably correlated with regional estimates of racial bias, but not with regional estimates of weight or gender bias. Riddle and Sinclair (2019) found that regional racial disparities in school-based disciplinary actions reliably correlated with regional estimates of racial bias, but not with regional estimates of sexuality bias.

² In this manuscript, we use the term ‘implicit’ to mean ‘indirect’. Thus, an ‘implicit measure’ assesses mental contents indirectly. We use the term ‘association’ to refer to one of the mental constructs assessed by implicit measures, but make no strong assumptions about the representational nature of the constructs assessed by implicit measures.

poor retest reliability (Gawronski et al., 2017; Nosek et al., 2005; Lai & Wilson, 2021). In contrast, in our investigation into the psychometrics of regional racial bias we found that measures of explicit and implicit racial bias both generally demonstrate poor retest reliability when aggregated to relatively small region levels (i.e., county), and only approach conventionally-acceptable levels of reliability at large region (i.e., state) levels. Similarly, whereas measures of explicit and implicit bias generally correlate modestly at the individual level, ranging $r = [0.14 - 0.35]$ (e.g., Bar-Anan & Nosek, 2014; Charlesworth & Banaji, 2019; Greenwald et al., 2009; Hofmann et al., 2005; Nosek et al., 2007; Oswald et al., 2013), we found that the correlation between explicit and implicit bias increased as the size of the regional unit increased, with state-level correlations as high as $r = 0.85$ (Hehman et al., 2019). This large state-level correlation suggests that regional measures of explicit and implicit bias may largely assess the same construct. However, interpretation of the state-level correlation is complicated by two related facts: aggregation increases the internal consistency of a measure by canceling out measurement error (Rushton et al., 1983), and internal consistency places an upper limit on how strongly two measures can correlate (Nunnally, 1970; Spearman, 1904). Consequently, an alternate interpretation of these high state-level correlations is that they simply reflect a known consequence of aggregation (i.e., canceling out measurement error to improve internal consistency). To test this possibility, we re-analyzed the aggregated state-level data with participants' locations randomly assigned.³ If aggregating individuals' responses on measures of

³ This approach is conceptually analogous to permutation tests used to evaluate machine learning algorithms (e.g., Olaja & Garriga, 2010). In a permutation test, one algorithm is trained to distinguish between two classes (e.g., pigs versus dogs) using a dataset in which the images are correctly labeled. Next, another algorithm is trained on the same dataset but with image labels randomly assigned (e.g., some pigs are labeled as dogs). In a subsequent testing phase, if the algorithm trained on the correctly-labeled images demonstrates a higher classification accuracy

explicit and implicit racial bias to region levels only cancels out measurement error, we should expect the observed large state-level explicit/implicit correlation to persist when participants are aggregated into states that do not reflect their true location. However, if location per se influences individuals' responses on measures of explicit and implicit racial bias, then the systematic influence of location will be lost when participants' locations are randomly assigned. We found support for the latter possibility, such that the state-level correlation between explicit and implicit bias was not reliably different from zero when participants' assigned state did not reflect their true location. Consequently, we interpret this pattern of results to indicate that measures of regional racial bias capture a meaningful construct that is geographically distributed.

INSERT FIGURE 1 ABOUT HERE

2. Theoretical Perspectives on Regional Intergroup Bias

Because psychology has traditionally conceptualized intergroup biases as attributes of people (e.g., Allport, 1954), bias is typically conceptualized to reflect a stable property of persons, and responses on measures of intergroup bias are interpreted to reflect the amount of bias a person carries with them from situation to situation. Consequently, regional differences in intergroup bias can be interpreted to reflect a greater proportion of biased people in some regions than in other regions. However, some theoretical perspectives also conceptualize bias as a property of contexts, and propose mechanisms by which bias can be created and perpetuated as a *macro-psychological* construct – one that exists beyond the mind of a single individual. We review these theories here, as well as other theoretical perspectives that are related to

rate than the algorithm trained on randomly-labeled images, then the correctly-trained algorithm is inferred to have identified a true class structure in the data.

conceptualizations of macropsychological phenomena more generally, inasmuch as they inform our understanding of regional intergroup bias.

2.1 Prejudice-in-Places

Building upon individual-level conceptualizations of bias, the Prejudice-in-places model (Murphy et al., 2018; Murphy & Walton, 2013) posits that intergroup biases can also exist within places. Specifically, a place can be characterized as prejudiced when it creates predictable and systematic inequalities in experience and outcomes because of people's social group memberships that afford advantages to people from some groups while disadvantaging people from other groups (Murphy & Walton, 2013). This distinction between prejudiced people versus places largely maps onto the sociological distinction between individual and institutional discrimination (Pincus, 1996).

The Prejudice-in-places model proposes a number of ways in which places can become prejudiced. Places can be prejudiced when practices, policies, or procedures disadvantage some groups relative to others, even when the practices seem neutral on their face (i.e., disparate impact). Similarly, places can be prejudiced when they rely on decision criteria that disadvantage members of some groups relative to others. For example, accredited American law schools' admission criteria typically weigh undergraduate GPA at about 40% and LSAT scores at about 60%; however, women have higher GPAs but lower LSAT scores than do men, on average (Wightman, 1998). Thus, criteria that would appear to apply to all groups equally instead systematically benefit one group over another.

Places can also be prejudiced when they impose greater psychological burden on some groups than others. For example, in contexts where certain groups are negatively stereotyped or numerically underrepresented, members of these groups can become physically and

psychologically vigilant for cues of belonging and inclusion – which, in turn, negatively impacts motivation and performance (Cheryan et al., 2009; Murphy et al., 2007; Walton & Cohen, 2007). Similarly, places can be prejudiced when members of some groups but not others must hide or change aspects of their identities in order to excel in the environment.

Taken together, the Prejudice-in-places model describes a variety of formal (e.g., policies) and informal (e.g., norms) ways in which places can be biased against members of some groups. Murphy and colleagues (2013, 2018) largely articulate the attributes of prejudiced places in terms of organizations, rather than in terms of geography, so the Prejudice-in-places model is not explicitly specified as a model of regional intergroup bias. Nevertheless, some tenets of the model can be readily scaled to regions. For example, at first glance, state budgeting priorities in Medicaid spending might seem to affect all state residents equally. However, racial minorities disproportionately rely on public health programs, so they are disproportionately impacted by policies that dictate Medicaid funding – and especially so in states with higher levels of regional racial bias (Leitner et al., 2018).

2.2 Bias of Crowds

Like the Prejudice-in-places model (Murphy et al., 2018; Murphy & Walton, 2013), the Bias of Crowds model (Payne et al., 2017a) conceptualizes intergroup bias as a property of contexts. However, the models differ in their focus on the constructs reflected in different measures of intergroup bias. Whereas the Prejudice-in-places model generally makes no strong distinction between the contributions of explicit versus implicit bias to contextual prejudice, the Bias of Crowds model proposes that implicit bias specifically reflects the influence of situational and environmental factors.

In support of this conceptualization of implicit intergroup bias as a property of contexts, the Bias of Crowds model offers three main pieces of evidence. First, implicit bias effects among individuals are robust: relative implicit preferences for one group over another consistently appear in the aggregate in theoretically-consistent directions (i.e., in favor of ingroups and/or high-status groups; Greenwald et al., 2009; Kurdi et al., 2019). Moreover, interventions to change individuals' implicit biases last only a short time, with average implicit biases quickly returning to their baseline levels (Lai et al., 2016). The consistency and durability of implicit biases suggest that implicit bias as a construct is relatively stable. The most prominently-used implicit measures (i.e., the IAT: Greenwald et al., 1998; and the AMP: Payne et al., 2005) also demonstrate good internal consistency, but routinely demonstrate relatively poor retest reliability (e.g., Bar-Anan & Nosek, 2014; Cunningham et al., 2001; Gawronski et al., 2017; Lai & Wilson, 2021). Taken together, this pattern of findings suggests the puzzling conclusion that implicit bias is stable in the aggregate but not stable within individuals. Second, and related, young children demonstrate implicit biases of similar magnitude to adults, which suggests that implicit bias is learned early and remains stable over the lifespan, but this conclusion is incongruous with observed low retest reliability (Gawronski et al., 2014; Lai & Wilson, 2021). Third, and perhaps most closely related to the present chapter, the average effect size for the relationship between individual-level implicit bias and behavioral outcomes is typically relatively small in magnitude (e.g., Greenwald et al., 2009; Forscher et al., 2019; Kurdi et al., 2019; Oswald et al., 2013), but regional implicit bias tends to exhibit relatively stronger relationships with behavioral outcomes (e.g., Hehman et al., 2018; Leitner et al., 2016a, b).

The Bias of Crowds model reconciles these seemingly-contradictory findings by concluding that implicit bias is not a stable property of individuals but, instead, reflects the

aggregate effect of individual fluctuations in context-dependent concept accessibility. Payne and colleagues attribute these contextual effects to a variety of factors, ranging from historical antecedents (Payne et al., 2019) to structural inequality (Vuletich & Payne, 2019) to biased mass media depictions (Weisbuch et al., 2009), segregated housing patterns, and norms and policies instituted in legal and education systems (Payne et al., 2017a). To the extent that structural features are relatively slow to change, we should expect regional aggregates of implicit bias to demonstrate higher retest reliability than individual intergroup bias – a prediction that is supported by previous research. Whereas the retest reliability of individual implicit bias generally ranges $r = 0.31 - 0.45$ across testing intervals of a few hours to a few months (Bar-Anan & Nosek, 2014; Cunningham et al., 2001; Gawronski et al., 2017), the year-over-year retest reliability of state-level implicit bias can be as high as $r = 0.76$ (Payne et al., 2017a; Hehman et al., 2019). Though not wholly discounting the influence of individual-level factors, Payne and colleagues conclude that implicit bias measures reflect cultural stereotypes and structural inequalities in the environment (Vuletich & Payne, 2019) and “...are better measures of situations or social environments than of persons” (Payne & Vuletich, 2018, p. 52). In turn, Payne and colleagues position the Bias of Crowds perspective as a bridge that connects traditional, individual-level theories of implicit bias to theories that conceptualize prejudice in terms of structures and institutions (Payne et al., 2017b).⁴

⁴ Though the IAT was developed to assess personal attitudes (Greenwald et al., 1998), some researchers have argued that responses on the IAT reflect “extrapersonal associations” (Olson & Fazio, 2004) or “environmental associations” (Karpinski & Hilton, 2001). To the extent that the IAT assesses constructs that originate – or exist entirely – outside of the individual, as Payne and colleagues suggest, then the evidence we review and the perspectives we propose in this chapter connecting regional biases to contextual factors are perhaps unsurprising. That said, concerns about extrapersonal or environmental associations have largely been articulated in the context of the IAT, but are not situated to explain regional bias effects obtained with other measures, either implicit or explicit. Consequently, these concerns highlight the importance of regional intergroup

2.3 Related Theoretical Perspectives

The Prejudice-in-places (Murphy et al., 2018; Murphy & Walton, 2013) and Bias of Crowds (Payne et al., 2017a) models both conceptualize intergroup bias as a property of places or contexts which, in turn, have implications specifically for regional intergroup bias as a construct. Other, related theoretical perspectives propose a variety of mechanisms explaining regional psychological phenomena more generally. For example, Rentfrow and colleagues (2008) propose a model of regional variation in psychological characteristics that could give rise to macro-psychological constructs. The model consists of three interrelated mechanisms: regional differences in the prevalence of psychological characteristics, regional differences in behaviors related to psychological characteristics, and geographic manifestations of psychological characteristics. Though their model is articulated to apply to psychological characteristics in general, Rentfrow and colleagues (2008) offer a specific example in the context of intergroup bias. If a relatively large proportion of biased people live in a region (e.g., because of selective migration), then biased behaviors (e.g., discrimination) should occur relatively more often in that region than in regions where smaller proportions of biased people live. People who live in a biased region but are not initially biased themselves may eventually become biased through repeated interactions with biased people (e.g., social influence) or by repeatedly observing biased behaviors (e.g., learning). Local clustering of biased people may eventually manifest regionally as outcomes (e.g., segregation) and institutions (e.g., biased laws) – and these biased outcomes and institutions can, in turn, perpetuate more biased behaviors and beliefs locally.

bias research that relies on measures other than the IAT – a point we elaborate on later in this chapter.

Other generalized theoretical perspectives also offer insight into the development or perpetuation of regional intergroup bias. For example, social impact theory (Latané, 1981) posits that local clustering of attitudes and beliefs can occur when individuals engage in repeated social interactions. Consequently, intergroup bias operationalized regionally may reflect the aggregate of the biases of the individuals that reside in a given region. However, regional bias may not reflect simply the aggregate of individual biases. Indeed, social psychology has been criticized for using the ‘group’ as shorthand for the average of individual phenomena (Hornsey, 2008), rather than treating groups and individuals as distinct levels of analysis. Other theoretical perspectives address this distinction between levels of analysis (e.g., Oishi & Graham, 2010; Pettigrew, 1997; Rentfrow et al., 2008) by proposing recursive relationships in which individual attitudes and beliefs become manifest in social structures (e.g., laws, institutions) that, in turn, influence the attitudes and beliefs of the individuals in that region (e.g., Ebert et al., 2020). Not only can such recursive relationships reinforce the individual biases that initially created the social structures, but they can also lead to emergent phenomena that are qualitatively distinct from the sum of their individual inputs (Smaldino, 2014).

3. The Puzzle of Explicit versus Implicit Intergroup Bias at the Region Level

Because intergroup bias has traditionally been conceptualized as a property of individuals (Allport, 1954; Murphy & Walton, 2013), explicit and implicit measures of intergroup bias were developed and validated to be used at the individual level based on theory articulated at the individual level. The task procedures of implicit measures were designed to constrain the influence of controlled processes (i.e., processes that are accessible to conscious introspection, can be deliberately initiated or terminated, require cognitive resources, and operate relatively slowly; Moors & DeHouwer, 2006) to a greater degree than do the task procedures of explicit

measures. Consequently, dual-process models of social cognition generally agree in interpreting individual-level responses on explicit and implicit bias measures to reflect distinct constructs (e.g., Cunningham et al., 2007; Fazio & Towles-Schwen, 1999; Gawronski & Bodenhausen, 2006; Strack & Deutsch, 2004). Supporting this distinction between individual-level explicit and implicit bias, a large body of evidence at the individual level shows that explicit and implicit measures independently predict behavior (Axt, Bar-Anan & Vianello, 2020), change at differing rates (Charlesworth & Banaji, 2019), and are sensitive to different information (Cao & Banaji, 2016). Additionally, measures of explicit and implicit bias correlate only modestly at the individual level. For example, among measures of individual racial bias, meta-analyses and other large-scale analyses converge on explicit/implicit correlations that range $r = [0.14 - 0.35]$ (Bar-Anan & Nosek, 2014; Charlesworth & Banaji, 2019; Greenwald et al., 2009; Hofmann et al., 2005; Nosek et al., 2007; Oswald et al., 2013). Similarly, among measures of individual sexuality bias, explicit/implicit correlations range $r = [0.17 - 0.43]$ (Charlesworth & Banaji, 2019; Greenwald et al., 2009; Nosek et al., 2007). That said, modest correlations between individual-level measures of explicit and implicit bias cannot be interpreted as incontrovertible evidence of distinct constructs because many implicit measures demonstrate poorer internal consistency than do explicit measures (see Gawronski et al., 2017), and internal consistency imposes an upper limit on the extent to which two measures can correlate (Nunnally, 1970; Spearman, 1904). However, the IAT (Greenwald et al., 1998) and the AMP (Payne et al., 2005) demonstrate conventionally-accepted levels of internal consistency (Gawronski & De Houwer, 2014; Gawronski et al., 2017; Lai & Wilson, 2021), and statistically accounting for measurement error produces only a medium-sized correlation between explicit and implicit measures (Cunningham et al., 2001).

The modest correlations between individual-level explicit and implicit bias are, at first glance, difficult to reconcile with recent findings from our own research on regional intergroup biases. When aggregated to region levels, correlations between estimates of explicit and implicit racial bias increase with the size of the regional unit (Hehman et al., 2019): explicit and implicit racial bias correlate as high as $r = 0.27$ at a relatively small (county) region level, $r = 0.77$ at larger (core-based statistical area) levels, and $r = 0.85$ at the largest (state) level. Similarly, measures of explicit and implicit sexuality bias correlate strongly at the state level, $r = 0.88$ (Ofosu et al., 2019). Whereas modest correlations between individual-level explicit and implicit bias can be interpreted to reflect either distinct constructs or differences in measurement error, substantial correlations between regional explicit and implicit bias offer relatively straightforward evidence that they largely assess the same construct.

Naturally, we cannot draw firm conclusions about the qualitative nature of regional intergroup bias based on observations in two domains (i.e., race, sexuality). Consequently, we report here meta-analyses of the correlations between explicit and implicit regional intergroup bias in Project Implicit (<https://implicit.harvard.edu>) data reflecting intergroup biases across 13 domains. We included responses from US-based participants who completed an IAT (Greenwald et al., 1998) that assessed the extent to which groups are associated with either evaluations (i.e., attitudes: Fazio et al., 1982), or attributes (i.e., stereotypes; Hamilton, 1981), as well as corresponding explicit measures. Race and Age data were collected between 2002 and 2020; Religion data were collected between 2004 and 2009; Sexuality data were collected in 2004 and between 2006 and 2020; Gender-science data were collected between 2003 and 2020; Gender-career data were collected between 2005 and 2020; and all other data were collected between

2004 and 2020. All data are available at <https://osf.io/y9hiq/>, and scripts for all analyses reported in this chapter are available at <https://osf.io/bsxc9/>.

As an index of implicit bias, we relied on the *D*-score (Greenwald et al., 2003), which quantifies in terms of response latency the extent to which one versus another social group facilitates responses to words related to evaluations (i.e., pleasant versus unpleasant) or attributes (e.g., harmless, dangerous). In the context of the attitude IATs, *D*-scores are coded such that higher values indicate a greater relative preference for the normatively higher status group over the normatively lower status group. In the context of the stereotype IATs, *D*-scores are coded such that higher values indicate stronger associations consistent with cultural stereotypes.

We computed indices of explicit bias to be analogous to the IAT *D*-score.⁵ For the attitude measures, we subtracted each participant's response on the single-item preference measure for the normatively lower status group from their single-item preference response for the normatively higher status group. For the stereotype measures, we subtracted each participant's response on the single-item attribute measure for the normatively lower status group from their single-item attribute measure for the normatively higher status group.

Based on each participant's self-reported location information, we aggregated *D*-scores and explicit measure difference scores into state-level estimates ($N = 51$, including Washington D.C.). Then, we calculated the correlation between state-level explicit and implicit bias for each

⁵ Previous regional intergroup bias research has operationalized regional explicit attitude bias in terms of feeling thermometers (Hehman et al., 2019; Ofofu et al., 2019), but there are no thermometers for the stereotype measures in the Project Implicit data. Consequently, in the present work we operationalized explicit bias in terms of single-item measures to be consistent across attitude and stereotype domains. Our attitude meta-analytic results were virtually identical when we used feeling thermometers instead of single-item explicit preference scores, which can be viewed at <https://osf.io/bsxc9/>.

intergroup domain, and meta-analyzed these correlations across domains, separately for attitude and stereotype measures, using the *R* package *metacor* (Balduzzi et al., 2019).

The results of these meta-analyses are depicted in Figures 2a and 2b. For the eight attitude measures, the meta-analytic correlation between state-level explicit and implicit intergroup bias is $r = 0.66$, which is a large effect according to Cohen's (1992) criteria (see also Lovakov & Agadullina, 2021), fixed effect 95% CI [0.60, 0.71], random effects 95% CI [0.32, 0.85]. For the five stereotype measures, the meta-analytic correlation between state-level explicit and implicit intergroup bias is $r = 0.33$, which is a medium effect, fixed effects 95% CI [0.21, 0.43], random effects 95% CI [0.18, 0.46].⁶

INSERT FIGURE 2 ABOUT HERE

The large meta-analytic estimate of the relationship between regional explicit and implicit attitude measures corroborates our previous findings in the domains of race and sexuality (Hehman et al., 2019; Oforu et al., 2019), and is far larger than individual-level correlations between the same measures reported previously (Bar-Anan & Nosek, 2014; Charlesworth & Banaji, 2019; Greenwald et al., 2009; Hofmann et al., 2005; Nosek et al., 2007; Oswald et al., 2013). At the same time, the magnitudes of regional explicit/implicit correlations are heterogeneous across attitude domains: correlations in the domains of sexuality and religion approach one, but correlations in the domains of disability and weight are not different from

⁶ We include confidence intervals for both fixed and random effects for the sake of completeness. However, the random effects confidence intervals should be interpreted cautiously because simulation studies indicate that relatively small numbers of clusters (as we have here) lead to biased estimates of standard errors (Mass & Hox, 2005).

zero. However, the meta-analytic estimate of the relationship between regional explicit and implicit stereotype measures is smaller than the estimate for regional attitude measures, and there is less heterogeneity in the magnitude of regional explicit/implicit correlations across stereotype domains than there is across attitude domains. That said, the regional stereotype correlation meta-analytic estimate is larger than individual-level correlations between explicit and implicit stereotype measures: very well-powered previous research at the individual level (Nosek et al., 2007) has revealed explicit/implicit correlations for the five stereotype measures we report here that range from $r = 0.15$ (Race-weapons) to $r = 0.22$ (Gender-science).

Several conclusions can readily be drawn from this pattern of results. First, the relatively small regional stereotype meta-analytic estimate, and the heterogeneity in regional attitude correlations, speak against the possibility that the large effect sizes that routinely emerge in regional analyses only reflect reduced measurement error as a function of aggregation (e.g., Connor & Evers, 2020). Aggregation reduces random error (Rushton et al., 1983), and random error places an upper limit on how strongly two measures can correlate (Nunnally, 1970; Spearman, 1904). Because aggregation raises the upper limit on how strongly two measures can correlate, aggregation may reveal larger (true) regional correlations but would not cause smaller (false) ones. Moreover, instead of aggregation uniformly inflating all correlations, the range of correlations reflected in our meta-analyses (which spans $[-0.03, 0.94]$) suggests that regional explicit and implicit bias measures reflect an underlying process (or processes) that varies across domains. Second, the large regional attitude meta-analytic correlation suggests that regional explicit and implicit attitude measures largely assess the same construct. However, this interpretation would seem to be at odds with theoretical perspectives that conceptualize

individual-level explicit and implicit bias as distinct constructs (e.g., Cunningham et al., 2007; Fazio & Towles-Schwen, 1999; Gawronski & Bodenhausen, 2006; Strack & Deutsch, 2004).

Two theoretical perspectives specifically focus on intergroup bias as a property of situations: the Prejudice-in-places model (Murphy et al., 2018; Murphy & Walton, 2013) and the Bias of Crowds model (Payne et al., 2017a). The Prejudice-in-places model does not distinguish between explicit versus implicit bias but, instead, focuses primarily on intentional (i.e., institutional) versus unintentional (i.e., structural) bias (see also Pincus, 1996). Consequently, the Prejudice-in-places model does not offer a perspective to interpret large correlations between regional measures of explicit and implicit bias. In contrast, the Bias of Crowds model (Payne et al., 2017a) focuses on implicit bias as a property of contexts and situations and, thus, would seem to be better positioned to explain large regional correlations. That said, the Bias of Crowds model is not articulated to account for explicit intergroup bias at either individual or region levels. Indeed, two of the three pillars upon which the Bias of Crowds model is built do not apply to explicit intergroup bias. Specifically, the poor retest reliability of individuals' implicit biases suggests that implicit bias is not a stable property of persons, but individuals' explicit biases demonstrate good retest reliability (e.g., Gawronski et al., 2017), which suggests that explicit biases are stable properties of persons. Similarly, the observed correspondence between the implicit biases of young children and adults suggest that implicit bias reflects the shared environments in which both children and adults reside, but children demonstrate higher explicit biases than do adults (e.g., Baron & Banaji, 2006), which suggests that explicit biases do not reflect shared environments. Thus, neither the Bias of Crowds model, nor any other existing theoretical perspective on intergroup bias that we are aware of, provides a mechanism to explain large correlations between regional explicit and implicit evaluative bias.

4. Toward a New Perspective on Regional Intergroup Bias

A large meta-analytic correlation between explicit and implicit attitude measures of regional intergroup bias suggests that both measures largely assess the same construct – at least in some intergroup domains – but existing intergroup bias theory cannot explain this pattern of results. Both theory and copious empirical evidence would seem to suggest that explicit and implicit measures assess distinct constructs at the individual level, but two issues stand in the way of us using information about individual-level intergroup bias to inform our understanding of regional intergroup bias.

First, we cannot use information about individual-level intergroup bias to draw conclusions about regional intergroup bias because levels of analysis are conceptually distinct. To assume a relationship observed at one level of analysis persists at other levels of analysis is to commit the ecological fallacy (Selvin, 1958; Simpson, 1951). Instead, evidence is necessary to determine whether a relationship observed at one level of analysis will persist at other levels. In the classic illustration of the ecological fallacy, American states with higher proportions of foreign-born residents had higher literacy rates (Robinson, 1950), suggesting the counter-intuitive conclusion that foreign-born people are more likely to be fluent in English than native-born Americans. When the same data were examined at the individual level, literacy rates were higher among natives than among foreign-born people: Immigrants are more likely to settle in states with greater educational and employment opportunities, which are also associated with greater English literacy. This illustration is not proof that a relationship at one level of analysis cannot persist at other levels; instead, it demonstrates the hazard of assuming correspondence across levels of analysis. Thus, we cannot assume that explicit and implicit measures of intergroup bias will correlate similarly at individual and region levels.

Second, we cannot draw firm conclusions about regional intergroup bias based on individual-level intergroup bias because of the disparity in internal consistency between individual-level explicit and implicit measures. Some implicit measures of intergroup bias, such as the IAT (Greenwald et al., 1998) and its variants and the AMP (Payne et al., 2005) routinely demonstrate conventionally-accepted levels of internal consistency, but other implicit measures, such as the evaluative priming task (Fazio et al., 1995), implicit relational assessment procedure (Barnes-Holmes et al., 2006), and extrinsic affective Simon task (De Houwer, 2003) do not (Gawronski & De Houwer, 2014; Gawronski et al., 2017; Lai & Wilson, 2021). Nevertheless, explicit measures of intergroup bias routinely demonstrate better internal consistency than do even the most internally-consistent implicit measures (Gawronski et al., 2017). Internal consistency imposes an upper limit on the extent to which two measures can correlate (Nunnally, 1970; Spearman, 1904), so observed modest correlations between individual-level measures of explicit and implicit bias cannot conclusively be interpreted as evidence that they reflect distinct constructs. Instead, differences in internal consistency at the individual level may be suppressing stronger correlations if explicit and implicit measures assess the same construct.

Though issues of internal consistency confound strong conclusions about the relationship between explicit and implicit bias at the individual level, internal consistency is less of a problem at region levels. When intergroup bias is operationalized as a regional construct, regional bias estimates reflect the aggregate of the responses of many individuals living in each region. The process of aggregation reduces unsystematic variance (e.g., random measurement error) because random error cancels itself out, thereby amplifying systematic variance (e.g., the construct of interest) and increasing the consistency of estimates (Rushton et al., 1983). Thus, differences in

internal consistency pose less threat to interpreting correlations between regional measures of explicit and implicit bias than at the individual level.

We propose that this principle of aggregation suggests a novel interpretation of explicit and implicit measures of intergroup bias when they are aggregated regionally – which, in turn, informs our understanding of the qualitative nature of regional intergroup bias. Not only does regional aggregation reduce measurement error (Rushton et al., 1983), but it also cancels out individual differences that are not shared by residents of a region, thereby amplifying psychological characteristics that are shared by residents of a region. In turn, this property of aggregation raises the question: To what extent do responses on explicit and implicit measures reflect individual differences versus shared characteristics? A wide variety of theoretical perspectives articulated at the individual level posit that responses on both types of intergroup bias measures reflect mental associations between groups (e.g., Black people; White people) and either evaluations (i.e., attitudes: good, bad) or attributes (i.e., stereotypes: e.g., athletic, intelligent) stored in memory (e.g., Cunningham et al., 2007; Fazio & Towles-Schwen, 1999; Gawronski & Bodenhausen, 2006; Strack & Deutsch, 2004). Some theoretical perspectives frame these associations as primarily reflecting the slow accumulation of direct, personal experience (e.g., Greenwald & Banaji, 1995; Wilson et al., 2000), whereas other perspectives conceptualize them to primarily reflect cultural learning (e.g., Karpinski & Hilton, 2001; Olson & Fazio, 2004).⁷ Though a variety of definitions of culture exist, several converge on the idea

⁷ These two interpretations are not necessarily mutually exclusive. Some perspectives view personal versus cultural interpretations of associations as a false dichotomy, e.g., “Implicit attitudes, as I see it, reflect traces of experiences within a culture that have become so integral a part of the individual's own mental and social makeup that it is artificial, if not patently odd, to separate such attitudes into ‘culture’ versus ‘self’ parts.” (Banaji, 2001, p. 139).

that culture reflects ideas, beliefs, or norms shared by people in a geographic region (e.g., Triandis, 2002; Varnum & Grossmann, 2017). Experiences may vary between persons (e.g., “idiosyncratic learning histories”, Gawronski & Bodenhausen, 2017, p. 269), but culture remains relatively more stable and constant within regions. Consequently, the process of aggregating individual measures of explicit and implicit bias into regional estimates should cancel out unshared individual experiences, leaving regional estimates of intergroup bias to primarily reflect the influence of shared culture (see also Charlesworth & Banaji, 2019).

Not only should aggregation cancel out unshared individual experiences, but it should also cancel out individual differences that are not directly related to intergroup bias but nevertheless influence responses on explicit and implicit measures. Many theoretical perspectives conceptualize explicit bias to reflect the influence of cognitive processes and other knowledge structures that constrain the expression of biased associations, such as social desirability and self-presentation concerns (e.g., Crandall & Eshleman, 2003; Roese & Jamieson, 1993), or egalitarian (e.g., Devine et al., 2002) and system justification (Jost et al., 2004) motivations. Similarly, formal mathematical modeling has identified a variety of processes besides associations that influence responses on implicit measures, such as inhibition, response biases, and task-simplification strategies (e.g., Conrey et al., 2005; Meissner & Rothermund, 2013; Nadarevic & Erdfelder, 2011; Stahl & Degner, 2007). Some of these processes and structures may vary between people (i.e., individual differences), but others may be common to people within a region (i.e., group processes). From this perspective, regional aggregation should cancel out the influence of cognitive processes and other knowledge structures that vary between people, but amplify the influence of processes and structures that are shared among people in a region.

In the context of intergroup bias, what cognitive processes or knowledge structures might be shared regionally? Though we know of no research that has investigated this question directly, it seems reasonable to posit that some motivations may vary regionally, perhaps according to local norms. For example, motivations to respond without prejudice (Plant & Devine, 1998) may be higher in regions that are relatively more egalitarian (Figure 3). To the extent that culture reflects ideas, beliefs, or norms shared by people in a geographic region, as cultural theorists have proposed (e.g., Triandis, 2002; Varnum & Grossmann, 2017), then motivations that are common to a meaningful proportion of the population in a region can reasonably be understood to reflect local culture.

INSERT FIGURE 3 ABOUT HERE

Synthesizing the steps we have outlined here leads to the following conclusion: Whereas individual-level measures of explicit and implicit intergroup bias reflect a combination of individual differences (e.g., Greenwald & Banaji, 1995; Wilson et al., 2000), cultural learning (e.g., Karpinski & Hilton, 2001; Olson & Fazio, 2004), and measurement error, aggregation to region levels attenuates the influence of individual differences and measurement error, leaving regional estimates of both measures to primarily reflect the influence of shared culture. Based on this reasoning, we propose that regional intergroup bias – measured either explicitly or implicitly – reflects a construct that is distinct from individual intergroup bias.

The perspective on regional intergroup bias that we have proposed here is a theoretical advancement because it highlights aggregation as substantively affecting the meaning of regional estimates of intergroup bias. That said, we are certainly not the first psychological scientists to

recognize aggregation as meaningful rather than mundane. For example, Kurdi and Banaji (2017) also note that regional aggregation cancels out both measurement error and individual difference in the context of implicit bias, and Garcia-Marques and colleagues (2017) make a similar observation in the context of explicit measures of stereotyping. Indeed, the latter even go a step further, and observe that "...aggregation and cancelation is the stuff cognition is made of" (Garcia-Marques et al., 2017, p. 266). Thus, we join an esteemed tradition of recognizing the implications of aggregation to provide novel insight into the qualitative nature of regional intergroup bias.

4.1 Implications and Caveats

4.1.1. Cultural Consensus and Social Priority

Though our meta-analysis indicated a large and reliable correlation between state-level measures of explicit and implicit attitudes, Figure 2 indicates a degree of heterogeneity in the magnitude of correlations across intergroup domains. Whereas state-level measures of explicit and implicit attitudes correlate strongly in the domains of sexuality, religion, race, and skin tone, state-level correlations do not differ from zero in the domains of weight and disability. Based on the perspective on regional intergroup bias we have proposed here, we should expect large correlations between regional aggregates of explicit and implicit measures when shared culture influences processes that contribute to responses on both explicit and implicit measures (e.g., associations). At the same time, we should expect small (or no) correlations between regional aggregates of explicit and implicit measures either when shared culture influences processes that contribute to responses on only one type of measure (e.g., self-presentation concerns), or when there is no cultural consensus on a topic. We speculate here about one mechanism that might underlie this pattern of correlations.

Inspired by the work of Charlesworth and Banaji (2019), we begin by proposing that variation in regional explicit/implicit correlations reflect different levels of cultural consensus as a function of social priority.⁸ Charlesworth and Banaji (2019) found that implicit biases in the domains of sexuality, race, and skin tone changed over time to become more egalitarian, but implicit biases in the domains of disability, age, and weight did not. They interpreted this pattern of results to reflect differences in the extent to which some topics versus others are prioritized in public discourse. In support of this interpretation, they presented data on the relative frequency of Google searches related to these different intergroup domains, with the idea that people will search more frequently for terms related to concepts that are prioritized in public discourse. Charlesworth and Banaji (2019) found that terms related to sexuality and race were searched far more frequently than were terms related to disability, age, and weight. These results suggest that the higher social priority of issues related to sexuality and race leads to increased cultural consensus in these intergroup domains relative to disability, age, and weight (cf. Levy & Banaji, 2002).

The pattern of change in individual-level intergroup biases over time observed by Charlesworth and Banaji (2019) maps perfectly onto our findings that regional explicit and implicit biases correlate more strongly in the domains of sexuality and race than in the domains of disability, age, and weight. Moreover, their proposed mechanism of social priority dovetails with the arguments we put forth here to suggest that the same cultural influences that contribute to stability in individual-level intergroup biases are also reflected in regional intergroup biases. Consequently, we synthesize their ideas with ours to propose that when strong cultural consensus

⁸ This conceptualization of cultural consensus also maps onto the construct of elaboration, which moderates explicit/implicit correlations at the individual level (Nosek, 2005).

exists regarding a given intergroup domain, and when this consensus influences processes that contribute to responses on both explicit and implicit measures, then the process of aggregation will amplify the effect of consensus in the estimates of both explicit and implicit bias and result in large correlations between the two measures.

We have proposed cultural consensus as a mechanism that would produce large correlations between regional measures of explicit and implicit intergroup bias. This mechanism also suggests two possible explanations for heterogeneity in regional evaluative correlations and for the relatively small regional stereotype meta-analytic estimate. One explanation is that there is relatively less public discourse about some intergroup domains than others; without public discourse to create cultural consensus, responses on both explicit and implicit measures will primarily reflect individual differences (e.g., individual experiences and beliefs), which will cancel out when aggregated to region levels and result in small (or null) correlations.

Another explanation for small (or null) regional explicit/implicit correlations is that public discourse in some domains creates cultural consensus that selectively manifests on processes that influences responses on only one type of measure (e.g., self-presentation concerns). When cultural influences are not reflected in both measures, or when different cultural influences are reflected in different measures, then the two measures will not correspond very much, if at all. In support of this latter explanation, the correlation between regional explicit and implicit weight bias does not differ from zero (Figure 2), which suggests that regional aggregates of the two measures reflect different influences. We are aware of only one published investigation into regional weight bias: Marini and colleagues (2013) found that nation-level implicit weight bias is higher in nations with more obese populations, but nation-level explicit weight bias is unrelated to population obesity. If our first explanation were correct – such that

there was little cultural consensus in the domain of weight – then responses on both explicit and implicit measures of weight bias would primarily reflect individual differences that cancel out with aggregation, and we would expect that neither form of bias would predict outcomes. In contrast to this prediction, Marini and colleagues' (2013) findings suggest that cultural consensus about weight manifests on processes that influence responses on implicit but not explicit measures. Nevertheless, future research will be necessary to tease these explanations apart, as well as to demonstrate more conclusively that public discourse contributes to regional intergroup biases.

Though the perspective we outline here suggests two explanations for small (or null) regional explicit/implicit correlations – lack of cultural consensus, or consensus that selectively manifests on processes that do not influence responses on both types of measure – two other explanations bear consideration. One possibility is that, rather than reflecting a lack of consensus within regions, small (or null) regional explicit/implicit correlations reflect considerable consensus across regions. For example, if explicit and/or implicit racial bias was uniformly strong (or uniformly weak) across all states in a nation, we would expect for the correlation between regional aggregates of the two measures to be low because variance is a necessary precondition for covariance. A cursory investigation into state-level variability in the Project Implicit explicit and implicit bias measures reported in our meta-analyses indicates that standard deviations in the attitude domain with the largest explicit/implicit correlation (Sexuality: explicit SD = 0.13, implicit SD = 0.05) are descriptively greater than in domain with the smallest explicit/implicit correlation (Disability: explicit SD = 0.04, implicit SD = 0.02). However, standard deviations in the stereotype domain with the largest explicit/implicit correlation (Gender-career: explicit SD = 0.08, implicit SD = 0.01) are descriptively less than in the

stereotype domain with the smallest explicit/explicit correlation (Asian-foreign: explicit $SD = 0.09$, implicit $SD = 0.06$). Given the mixed results of these relatively superficial comparisons, the question of variability (or lack thereof) in intergroup biases across intergroup domains merits further, more formalized investigation.

Another possibility is that small (or null) regional explicit/implicit correlations reflect procedural differences between explicit and implicit measures that affect their construct validity. For example, explicit measures of weight bias rely on the terms ‘thin’ and ‘fat’, whereas implicit measures of weight bias rely on silhouettes of people with varying body sizes. Stimuli necessarily influence the extent to which a measure assesses its intended construct, so it is possible that regional measures of explicit and implicit weight bias do not correlate strongly because the body shape silhouettes do not assess the same construct as do the category labels. Of course, this critique is potentially true of all measures and domains – including the ones with high explicit/implicit correspondence – and, thus, is not a parsimonious explanation for the observed small (or null) correlations. Nevertheless, future research should investigate these alternate explanations for variability in regional explicit/implicit correspondence.

4.1.2. Does the Regional Perspective Inform Individual-Level Perspectives on Implicit Bias?

Our perspective on regional intergroup bias is positioned to contribute to the ongoing debate inspired by Payne and colleagues’ Bias of Crowds model (Payne et al., 2017a) over the extent to which implicit bias is better conceptualized as a property of situations, a property of individuals, or both (e.g., Connor & Evers, 2020; Daumeyer et al., 2017; Gawronski & Bodenhausen, 2017; Kurdi & Banaji, 2017; Rae & Greenwald, 2017; Samayoa & Fazio, 2017). Importantly, however, this debate is centered on implicit bias as an individual-level construct,

whereas here we focus on intergroup bias as a regional construct. Effects at one level of analysis do not necessarily generalize to other levels (Selvin, 1958; Simpson, 1951), and we have proposed a rationale based on aggregation that conceptualizes regional intergroup bias and individual intergroup bias to reflect distinct constructs. Thus, one implication of our work is that conclusions about the nature of individual-level intergroup bias should not be drawn from evidence based on regional investigations. Thus, we expressly refrain from drawing such conclusions here, and we caution other researchers against doing so, as well.

The perspective we propose here also has implications for perspectives on individual-level bias that rely on the principles of aggregation. For example, Connor and Evers (2020) use real and simulated data to show that Payne and colleagues' (2017a) arguments in favor of conceptualizing implicit bias as a property of contexts can be explained in terms of measurement error that is reduced through aggregation. Though Connor and Evers (2020) conclude that "implicit bias is best understood as an individual-level construct measured with substantial error" (p. 1329), they also recognize that implicit biases can vary regionally, in the form of individuals with similar levels of bias clustering in groups. From this perspective, they interpret large correlations between regional intergroup bias and regional outcomes (e.g., Hehman et al., 2019) as "artifact(s) of aggregation plus some level of systematic group-level variation." However, regional implicit bias is not just a less noisy version of individual-level implicit bias. Instead, regional implicit bias reflects shared cultural influences to a greater degree, and individual differences to a lesser degree, than does individual-level implicit bias. Consequently, in our view individual and regional intergroup bias should not be conceptualized as interchangeable versions of one another but, instead, be recognized and treated as qualitative distinct constructs, to be used appropriately when one or the other is more relevant to a particular research question.

Our critique would also seem to apply to suggestions that aggregating the responses of single individuals over multiple measurement occasions should have similar effects as aggregating the responses of multiple individuals who share a geographic region (e.g., Connor & Evers, 2020; Kurdi & Banaji, 2017; Machery, 2017). This suggestion is rooted in the idea that aggregation would cancel out measurement error in both cases, resulting in more statistically precise estimates of bias. Based on this perspective, Kurdi and Banaji (2017) hypothesized that the large difference in predictive power between individual-level and regional intergroup bias that Payne and colleagues (2017a) observed would shrink considerably or disappear altogether if individual-level studies relied on aggregated measures within individuals or otherwise accounted for measurement error. However, there are several reasons we should not expect aggregates within individuals to be conceptually similar to aggregates within regions. First, individual aggregates would amplify the influence of whatever is stable within individuals, but stable individual differences are canceled out in regional aggregates if they are not also shared among individuals within regions (e.g., idiosyncratic learning). Second, individual aggregates would amplify the influence of whatever is stable in the individual's measurement context. For example, to the extent that situational cues affect concept accessibility (Payne et al., 2017a), a person who completes an intergroup bias measure once on their multicultural college campus and again in their homogenous hometown should demonstrate less intergroup bias on the first than the second measurement. In contrast, a person who completes multiple intergroup bias measures within the same testing context should demonstrate relatively more stability across measurements. Regional aggregates necessarily amplify the influence of whatever is stable within a region, so the extent to which regional and individual aggregates both reflect contextual influences depends on the stability of the individual's context. Third, the mere fact of completing

the same intergroup bias measure multiple times can affect responses. This phenomenon of practice effects has been shown in the context of the IAT, such that people generally demonstrate the strongest bias on their first IAT compared to subsequent administrations (Greenwald et al., 2003). Aggregating responses within individuals would therefore amplify the influence of practice effects, whereas aggregating responses within regions would not. Thus, even if an intergroup bias measure primarily reflects contextual influences rather than individual differences (e.g., Payne & Hannay, 2021), individual aggregates would reflect contextual influences plus practice effects whereas regional aggregates would reflect contextual influences. Taken together, aggregates within individuals should reflect a different collection of influences (i.e., stable properties of individuals, stable properties of individuals' contexts, practice effects) than aggregates within regions (i.e., stable properties of regions), so they should not necessarily be expected to predict the same outcomes. Interestingly, Payne and Hannay (2021) come to the same conclusion as we do here about the conceptual distinction between individual and regional aggregates, but from a different rationale. To the extent that implicit bias is not a stable property of individuals, as they argue, then there should be little consistency in bias within individuals for aggregation to amplify, and aggregated repeated measures of individuals should not provide more predictive power than single measures. Indeed, Hannay and Payne (2021) found that aggregating the responses of individuals' implicit biases measured repeatedly over time provides only minimal additional predictive validity compared to a single measure. That said, Carpenter and colleagues (2021) employed largely the same paradigm as did Hannay and Payne (2021) but also statistically accounted for measurement error; they found the opposite pattern of results, such that aggregating individuals' repeated responses significantly improved predictive validity over a single measurement. So the jury is still out on the benefits to predictive validity of

aggregating implicit biases within individuals. Nevertheless, both of these investigations aggregated individuals' responses over only a handful of measurement occasions (range: 2-6), which is orders of magnitude smaller than the number of responses typically reflected in regional intergroup bias aggregates. Future research should continue to investigate intergroup biases aggregated within individuals versus regions, but feasibility may ultimately prohibit a true apples-to-apples comparison between individual and regional aggregates of intergroup bias.

4.1.3 Why the Geographic Unit of Analysis Matters

Our perspective on the effects of aggregation would seem to have direct implications for how we understand the role of the geographic unit of analysis in regional intergroup bias. For example, we observed previously that explicit/implicit correlations increase as regional units increase in size (Hehman et al., 2019). Usually (but not always), larger regions reflect the responses of more people. Consequently, when relatively few responses are aggregated into regional estimates – whether it be a geographically small region, such as postal code, or a sparsely-populated large region, such as Wyoming – idiosyncratic experiences and individual differences necessarily exert relatively more influence on the estimate of regional bias compared to shared cultural influences. Different individual differences influence responses on explicit versus implicit measures (e.g., self-presentation concerns versus task-set simplification strategies, respectively), which – along with the presence of relatively more measurement error in implicit measures – attenuates correlations between the two measures when regional estimates reflect relatively few responses. In contrast, when relatively more responses are aggregated into regional estimates – whether it be a densely-populated small region, such as Singapore, or a geographically large region, such as a nation – idiosyncratic experiences, individual differences,

and measurement error are canceled out to a greater degree relative to shared cultural influences, which results in larger correlations between measures.

Taken together, our predictions and observations suggest that there is not a “right” geographic unit or sample size for regional intergroup bias research. Because aggregation reduces statistical noise, we should expect the precision of regional intergroup bias estimates to increase with the number of observations per regional unit. However, the precision of an estimate also depends on the strength of the underlying signal: In general, we should expect regional estimates to be more precise in domains with greater versus less cultural consensus. Consequently, a small sample may produce precise estimates in a domain that is characterized by a high degree of cultural consensus, but even a very large sample may not produce precise estimates of regional intergroup bias in a domain that is characterized by a low degree of cultural consensus. Thus, we should expect regional estimates of intergroup bias to have the greatest precision when two criteria are met:

- when there is greater social consensus in the domain / topic; and
- when there are more observations to amplify social consensus and minimize measurement error and individual differences.

We acknowledge that these two criteria are very subjective: How much social consensus is sufficient? How many observations are sufficient? We hope that future research will help us to articulate these guidelines more concretely. For example, the published literature on regional intergroup bias will soon reach the point that it is ready to be meta-analyzed. A meta-analysis may identify an average effect size of regional intergroup bias, and moderators such as intergroup domain, geographic unit, or others that increase or decrease this average effect size –

which in turn would provide a degree of objective precision to our currently subjective criteria. Such a meta-analysis would seem to be a very fruitful direction for future research.

5. Regional Intergroup Bias: Next Steps Forward

The perspectives proposed here about aggregation changing the qualitative nature of measures of intergroup bias suggest new directions for intergroup bias research. With new opportunities come new questions to be answered and new problems to be solved. In this section we discuss what are, in our view, some of the important challenges that regional intergroup bias research will need to address in moving forward – and, when possible, we propose solutions. This section is generally organized according to topics related to: predicting outcomes and inferring causality; constructs, measures, and data sources; and levels of analysis.

5.1 Predicting Outcomes and Inferring Causality

5.1.1. Converging Evidence across Explicit and Implicit Intergroup Bias Measures

In domains where explicit and implicit measures of regional intergroup bias correlate strongly, as revealed by our meta-analyses, the two types of measures largely assess the same construct and therefore should be expected to predict the same outcomes. Indeed, this pattern of results can be observed in the published literature. Among sixteen investigations that modeled explicit and implicit bias separately in the domains of sexuality or race – which correlate at $r_s = 0.94, 0.86$, respectively, in our meta-analysis – each form of bias predicted the same outcome of interest in thirteen (81%) of them (Chin et al., 2020; Ekstrom et al., 2021; Hehman et al., 2019; Johnson & Chopik, 2019; Leitner et al., 2018; Ofosu et al., 2019; O'Shea et al. 2019; Pearman, 2021; Rae et al., 2015; Stelter et al., 2021)⁹. In contrast, all three of the investigations that

⁹ Chin and colleagues (2020) and Hehman and colleagues (2019) each report two investigations in which explicit and implicit bias are modeled separately, and in these cases both forms of bias predicted the same outcomes. Johnson & Chopik (2019) report four analyses in which explicit

modeled both forms of bias separately in domains where our meta-analyses revealed small correlations – weight ($r = 0.06$) and Asian-foreign stereotypes ($r = 0.16$) – regional explicit and implicit bias did not predict the same outcomes (Devos & Sadler, 2019; Devos et al., 2019; Marini et al., 2013). Thus, and perhaps unsurprisingly, regional explicit/implicit correspondence appears to determine the extent to which implicit versus explicit measures of regional bias predict the same outcomes.

When regional measures of explicit and implicit regional intergroup bias are modeled separately, a relatively consistent pattern of results emerges, such that both measures generally predict the same outcomes in domains with high regional explicit/implicit correspondence, but consistently do not predict the same outcomes in domains with low regional explicit/implicit correspondence. However, a more mixed pattern of results emerges when regional explicit and implicit bias are modeled together.¹⁰ Of nine investigations that model both forms of bias together in the domains of race or sexuality, only explicit bias predicted the outcome of interest in five of them (Leitner et al., 2016a; Stelter et al., 2021; Orchard & Price 2017)¹¹, only implicit bias predicted the outcome of interest in two of them (Herman et al. 2018; Miller et al., 2016), and both forms of bias simultaneously predicted the outcome of interest in one of them (Leitner

and implicit bias are modeled separately: in two cases both forms of bias predicted the same outcomes, and in two cases only explicit bias predicted outcomes.

¹⁰ Riddle and Sinclair (2019) do not fit cleanly into either of these summaries because they modeled explicit and implicit bias both separately and together to predict five different outcomes in two separate waves of data. In the analyses where explicit and implicit bias were modeled separately, explicit bias predicted nine of 10 outcomes and implicit bias predicted four of 10 outcomes. In the analyses where explicit and implicit bias were modeled together, explicit bias predicted seven of 10 outcomes and implicit bias predicted no outcomes.

¹¹ Leitner and colleagues (2016a) and Orchard and Price (2017) each report two investigations in which explicit and implicit bias are modeled together, and in these cases only explicit bias predicted the outcomes of interest.

et al., 2016b).¹² This relatively heterogeneous pattern of results may reflect the fact that both forms of bias largely provide redundant information about the outcome (i.e., are multicollinear), which can lead to Type II error (Lavery et al., 2019). Given that correlations between regional explicit and implicit intergroup bias depends both on the domain (as our meta-analyses indicate) and on the unit of analysis (Hehman et al., 2019), future regional intergroup bias research should include standard multicollinearity checks (e.g., variance inflation factor) or latent factor modeling in order to safeguard the predictive validity of their statistical models.

5.1.2 Establishing Causal Evidence

Any mature science will ultimately want to identify causal relationships. However, to date, almost all regional intergroup bias research has relied on correlational analytic methods that provide limited evidence for causality. That such large-scale studies employ correlational analyses is perhaps unsurprising, given available observational data and the logistical and ethical challenges associated with conducting experiments at region levels. Nevertheless, correlational approaches necessarily constrain conclusions that can be drawn about causality, so regional intergroup bias researchers may seek alternate approaches to infer causal relationships.

INSERT FIGURE 4 ABOUT HERE

When regional data has been collected over time, researchers can apply longitudinal designs to reveal trends over time (Figure 4) that can provide relatively strong evidence for

¹² We are aware of only one investigation that included both forms of bias in the same model in a domain other than race or sexuality: when both forms of bias are modeled together in the domain of Gender-science stereotypes ($r = 0.29$), only implicit gender stereotypes predict the outcome of interest (Nosek et al., 2009).

causality. We applied such an approach to examine the extent to which government legislation influenced the attitudes of citizens in the context of same-sex marriage (Ofosu et al., 2019). Our longitudinal design takes advantage of the quasi-experimental manner in which marriage-equality legalization unfolded on a state-by-state basis across the United States over the course of several decades. Comparing trends in bias before versus after legalization, we found that the rate at which attitudes towards gay people were becoming more egalitarian increased after relative to before legalization, regardless of when same-sex marriage was legalized in any given state.

Despite the correlational nature of these data, this multiple-groups, staggered-treatment, pre-post design provides relatively strong evidence of a causal relationship (i.e., that legislation influenced bias). However, this approach is still limited in that it does not speak to the reverse relationship, that changes in anti-gay bias also influenced legislation, nor can it examine reciprocal relationships between legislation and bias. Other approaches allow for greater causal inference from correlational data, and future regional intergroup bias research might rely on auto-regressive latent growth curve modeling (Berry & Willoughby, 2017), panel vector autoregression modeling (e.g., Götz et al., 2021), other longitudinal analyses from the cross-lag family of models (e.g., Grimm et al., 2021; Usami et al., 2019), or instrumental variable estimation (Angrist & Pischke, 2008).

5.1.3 Cultural Influence and Cultural Relevance

To the extent that regional intergroup bias largely reflects shared culture, we can speculate about factors that moderate the relationship between regional intergroup biases and regional outcomes. For example, if different groups of people vary in their influence on regional culture, we should expect that the biases of groups of people within a region will predict outcomes as a function of that group's influence on regional culture, and also as a function of the

relevance to the outcome. For example, we found that disproportionate use of lethal force against Black people by police is related to the regional biases of White but not Black people (Hehman et al., 2018). This pattern of results may suggest that White people exert more influence on police officers' behavior, or on the structures and institutions that influence police officers' behavior, than do Black people – or, conversely, that police officers' behavior more strongly influences the biases of White people than of Black people. On the other side of this coin, Price and colleagues (2021a, 2021b) found that psychotherapy interventions are less effective for Black but not White youth in more racist regions, and are less effective for girls but not boys in more sexist regions. To the extent that Black people are more affected by racism than are White people, and girls are more affected by sexism than are boys, this pattern of results would seem to indicate that the predictive validity of regional intergroup biases depends on the relevance of the bias to the outcome. As regional intergroup bias continues to grow as an area of research, and a nomological network of associated effects grows with it, we hope that theory will develop in its ability to better predict when and for whom regional biases are related to outcomes.

5.1.4 Intergroup Bias as a Cause versus Consequence

In laboratory-based, experimental intergroup bias research, independent and dependent variables can easily be distinguished from one another: the independent variable is manipulated (e.g., assigning participants to different bias-reduction interventions), and the dependent variable is measured (e.g., scores on an IAT). However, all variables are usually measured in regional intergroup bias research, so the line between bias and the consequence of bias can become blurred. For example, if regional intergroup bias as a construct primarily reflects shared culture, as we argue here, then regional bias manifest as prejudice would be a consequence of culture. However, to the extent that regional intergroup bias is related to discriminatory behaviors that

happen within regions (e.g., disparate police shootings), then bias either causes the discriminatory behaviors or both bias and the behaviors reflect an unobserved third variable (as is possible in all correlational research).

In some cases, temporal precedence can provide suggestive evidence for one causal pathway over the reverse. For example, estimates of regional racial bias from the early 21st century are higher in areas of America that had relatively higher enslaved populations before the Civil War in the mid-19th century (Payne et al., 2019), which suggests that historical structures influence present-day intergroup biases rather than the impossibility that present-day intergroup biases influenced historical structures. But in many other investigations of regional intergroup bias, evidence for temporal precedence or other indices of causality is less straightforward.

With this limitation in mind, we sought to identify characteristics of the environment that are related to regional intergroup bias but are logically unlikely to be direct consequences of bias (Hehman et al., 2020). We began by compiling a very large dataset of environmental attributes ($N > 800$) based on administrative data reflecting population demographics, health and healthcare metrics, topographical features, weather, temperature, and crime. We excluded variables that could plausibly be expected to be the direct or indirect by-products of intergroup biases, as well as variables whose temporal precedence to bias could not be clearly established. Then, we applied modern model-selection techniques to the administrative dataset to develop parsimonious predictive models of regional explicit and implicit measures of race-, religious-, sexuality-, age-, and health-based intergroup biases. We found that human features of the environment (e.g., commute length; availability of mental health providers) and events in the environment (e.g., rates of premature deaths; rates of obesity) consistently predicted biases (Figure 5). Moreover, these data-driven statistical models of environmental features predicted up

to 62% of variance in regional intergroup biases, and predicted significantly more variance in regional biases than basic regional demographics alone. Though this approach does not position us to make strong claims about causality (e.g., that regions are more biased because of poor urban planning that results in longer commute lengths), it nevertheless brings us closer to inferring causality by identifying environmental characteristics that are related to regional intergroup bias but are unlikely to be direct consequences of bias.

INSERT FIGURE 5 ABOUT HERE

Like in all correlational research, distinguishing causes from consequences will continue to be an issue in regional intergroup bias research. Moreover, in many cases the causal relationship between bias and consequences of bias is likely bi-directional, such that regional prejudice causes regional discrimination and regional discrimination reinforces regional prejudice. Nevertheless, our findings help to pave the way for more precise theorizing about causal mechanisms underlying regional intergroup bias.

5.2 Advancing Beyond Specific Constructs, Measures, and Data Sources

The majority of research to date examining regional intergroup bias has been based on data from the same source: Project Implicit. Project Implicit has existed since 2002 as a demonstration website, measuring visitors' attitudes and stereotypes across a wide variety of domains. Without question, Project Implicit has transformed psychological science by providing researchers with access to samples that are orders of magnitude larger than any lab-based study could provide. Additionally, the Project Implicit sample is far more diverse than a typical sample of university undergraduates on the dimensions of race/ethnicity, age, political orientation, socio-

economic status, and educational background. Because the demographics collected by Project Implicit include geolocation information, data can be readily aggregated into regional estimates.

Not only has most regional intergroup bias research to date relied on data from Project Implicit, but – as our literature review indicates – it has also largely focused on Black/White racial bias in American samples. This focus on one population, one construct, and one data source limits the generalizability of findings. Fortunately, other large-scale, geographically-rich datasets exist. For example, in the United States, the American National Election Studies (<https://electionstudies.org>), the General Social Survey (<https://gss.norc.org/>), the Public Religion Research Institute (<https://www.prrri.org>), and the Pew Charitable Trust (<https://www.pewresearch.org>) routinely collect measures of intergroup bias along with geographic information. Similar sources of information are available in the World Values Survey (<https://www.worldvaluessurvey.org>), which spans nearly 100 countries, as well as in the European Social Survey (<https://www.europeansocialsurvey.org>) and the European Values Study (<https://europeanvaluesstudy.eu>). For the purposes of regional intergroup bias research, these data sets are somewhat limited because they generally contain cruder measures of intergroup biases (e.g., single-item), fewer observations per geographic unit, and coarser spatial resolution (i.e., reflecting larger rather than smaller areas) than are available from Project Implicit. That said, these alternative data sources often reflect representative samples, and can be used in conjunction with more psychologically-focused data in order to provide converging evidence from multiple, independent sources. In Ofosu et al. (2019), we adopted this approach by first demonstrating a relationship between same-sex marriage legalization and state-level sexuality biases using data from Project Implicit, and then replicated our finding using data from the American National Election Studies.

Though the majority of regional intergroup bias research has relied on data from Project Implicit, several studies have incorporated data from other sources and may provide fruitful templates for future research. For example, in Esposito and Calanchini (2021) we examined person-environment fit in the context of sexuality: we used Project Implicit visitors' responses to index sexuality attitudes, and operationalized regional gay friendliness using nationally-representative data from the Public Religion Research Institute that assessed support for policies related to sexual minorities (e.g., same-sex marriage legalization). Miller and colleagues (2016) went a step further and collected their own data to examine the relationship between community-level biases and the well-being of people living with HIV. They recruited people living with HIV through clinics across New England, and recruiting community members through random-digit telephone dialing. In contrast to Miller and colleagues' (2016) relatively more traditional (and resource-intensive) approach to collecting regional data, Kunst and colleagues (2017) studied relationships among regional intergroup bias, violence, and economic inequality using crowdsourcing platforms. Specifically, they collected data through Amazon's Mechanical Turk, and set a threshold of 100 responses per state, which they achieved in 30 states.

Another approach, advocated by Rentfrow and colleagues (Rentfrow, 2010; Rentfrow & Jokela, 2016), is to combine data collected from different laboratories, and operationalize collection site as the unit of analysis. Vuletich and Payne (2019) adopted a version of this approach, using data that were originally collected by Lai and colleagues (2016) as part of a multi-labs effort across 17 American universities, which were initially analyzed with the individual participant as the unit of analysis. Vuletich and Payne (2019) re-analyzed these data, aggregating individual responses into university-level estimates, to examine the relationship between university structural features (e.g., faculty diversity; presence of confederate

monuments) and regional racial bias. Hatzenbuehler and colleagues (2021) also relied on this approach: they re-analyzed neuroimaging data from the Adolescent Brain Cognitive Development study, aggregated to the level of the collection site, to examine the relationship between children's neurodevelopment and regional racial bias.

A handful of studies have formalized this re-analysis approach to regional intergroup bias research by conducting spatial meta-analyses (Johnson et al., 2017). A spatial meta-analysis is similar to a traditional meta-analysis, but also includes geolocation information for each effect reflected in the meta-analysis. Price and colleagues (2021a; see also Price et al 2021b) provide an exemplary demonstration of using spatial meta-analysis to re-analyze laboratory-based data at the region level. To investigate the extent to which regional intergroup bias moderated the effects of psychotherapy interventions among youth, they drew more than 2,000 effect sizes reflecting psychotherapy efficacy from existing literature, and also recorded the racial demographics of each sample (i.e., majority Black; majority White) and the location (i.e., U.S. state) of the laboratory that collected each sample. Then, to operationalize state-level racism, they standardized and aggregated a composite of items related to Black/White intergroup relations from three sources: the American National Election Studies, the General Social Survey, and Project Implicit. Their meta-analysis revealed that state-level racism and racial demographics interacted to moderate the efficacy of psychotherapy. Given that intergroup bias has been studied in laboratories for decades and has produced a vast literature, spatial meta-analysis would seem to be an especially promising approach for future research.

As multi-lab efforts become more common, scientific norms continue to move towards open science practices, and infrastructure like the Psychological Science Accelerator (e.g., Moshontz et al., 2018) continue to develop, these kinds of re-analyses are positioned to advance

research on regional intergroup bias. As an added benefit, re-analyses approaches such as these help to ameliorate concerns about the quality of regional data, which are often collected online: Whereas researchers cannot control (or even know) the conditions under which online participants complete experimental tasks (e.g., in a busy coffee shop; while watching TV), laboratory data is often collected under very controlled conditions, which may reduce noise and, thus, increase statistical power.

5.2.1 American National Election Studies and the Affect Misattribution Procedure

In order for regional intergroup bias research to continue to advance, it cannot be based so heavily on a single data source (i.e., Project Implicit). Such disproportionate dependence on one data source – albeit a very large, psychologically comprehensive, and relatively diverse data source – limits the external validity of existing regional intergroup bias research. Moreover, in relying so heavily on data from Project Implicit, extant regional research on implicit bias specifically can also be characterized as the study of the IAT: We are aware of only one example of regional intergroup bias research that focused on implicit bias but did not use data from Project Implicit (Miller et al., 2016), and even that study used an IAT.

Other implicit measures, such as sequential priming tasks (e.g., Fazio et al., 1995; Payne et al., 2005), aim to assess the same construct as the IAT: mental associations. However, the IAT differs both procedurally and conceptually from sequential priming tasks, and task demands necessarily dictate which cognitive processes drive responses (Payne et al., 2008). For example, the IAT relies on category labels whereas sequential priming tasks generally do not, and category salience moderates implicit bias at the individual level (Olson & Fazio, 2003). Thus, the generalizability of regional intergroup bias findings to other operationalizations of implicit bias besides the IAT is an open question.

Here, we begin to investigate regional intergroup bias operationalized in terms of other implicit measures. We relied on nationally-representative data collected by the 2008 American National Election Studies (ANES) Time Series Study. As an implicit measure of intergroup bias, this wave of the ANES included an affect misattribution procedure (AMP: Payne et al., 2005) assessing evaluations of Black and White people. As an explicit measure of intergroup bias, it included feeling thermometers assessing evaluations of Black and White people.

The AMP is typically scored in terms of proportion of “more pleasant” versus “less pleasant” responses to neutral Chinese pictograph target images that follow prime pictures of White people and, separately, Black people. For each ANES respondent, we subtracted the proportion of “more pleasant” Black responses from the proportion of “more pleasant” White responses to create a difference score that is analogous to the IAT *D* score (Greenwald et al., 2003), such that higher scores reflect more positive evaluations of White people relative to Black people. We subtracted each respondent’s response on the Black feeling thermometer from their response on the White feeling thermometer.

In this sample of $N = 2040$ ANES respondents, individual-level implicit racial bias ($M = 0.13$, $SD = 0.24$), correlated with explicit racial bias ($M = 1.17$, $SD = 19.72$) at $r = 0.31$, 95% CI [0.27, 0.34], which closely aligns with the meta-analytic correlation between the AMP and explicit attitude measures ($r = 0.30$, 95% CI [0.25, 0.35]) reported by Cameron et al. (2012). Additionally, based on each participant’s self-reported location information, we aggregated AMP difference scores and feeling thermometer difference scores into state-level estimates ($N = 34$). At the state level, with 5,000 bias-corrected bootstraps we estimated the relationship between implicit ($M = 0.10$, $SD = 0.04$) and explicit racial bias ($M = 0.51$, $SD = 5.95$) as $r = 0.53$, 95% CI [0.24, 0.74]. This replicates the pattern of results we have observed in our own research

(Hehman et al., 2019; Ofofu et al., 2019), in that the correlation between explicit and implicit bias increases as the regional unit of analysis increases in size. In turn, these results suggest that the benefits of aggregation to explicit/implicit correspondence are not limited to a single implicit measure (i.e., the IAT) or to a single participant sample (i.e., Project Implicit visitors). However, the correlation between explicit and implicit racial bias in the ANES sample is not reliably different between individual ($r = 0.31$) versus state levels ($r = 0.53$) of analysis, based on a two-tailed Fisher r -to- z transformation, $z = -1.49$, $p = .136$, 95% CI of difference $[-0.43, 0.08]$. That said, this ANES sample is far smaller than the Project Implicit race sample, both in terms of number of participants ($Ns = 2,040$ versus 3,179,403, respectively) and number of states ($Ns = 34$ versus 51, respectively). In light of these differences, we are hesitant to over-interpret this apparent null difference between levels of analysis in the ANES sample. Nevertheless, these findings represent an important initial investigation into regional intergroup bias in data from sources other than Project Implicit, and from measures other than the IAT.

5.2.2 Project Implicit versus the American National Election Studies

Given that such a large proportion of the published regional intergroup bias literature reflects a single participant sample and a single measure (i.e., Project Implicit and the IAT), we also examined the extent to which regional intergroup bias corresponds across participant samples and measures. We reported initial evidence on this point in Ofofu et al. (2019), with state-level estimates of sexuality feeling thermometers from Project Implicit and the ANES correlating at $r = 0.75$. Here, we extend this investigation to implicit measures by correlating state-level racial IAT estimates from Project Implicit with state-level AMP estimates from the ANES. The ANES data were collected in 2008, so we restricted the Project Implicit data to responses collected in 2008.

Among the 34 states common to both datasets, state-level AMP and IAT correlated with 5,000 bias-corrected bootstraps at $r = 0.33$, 95% CI [0.05, 0.56]. This correlation aligns with an individual-level correlation between the AMP and IAT of $r = 0.30$ (Bar-Anan & Nosek, 2014), but is far smaller than the $r = .75$ correlation between state-level explicit bias measures we observed previously in the same two data sources (Ofosu et al., 2019). However, the feeling thermometers used by Project Implicit and the ANES are structurally identical, whereas the IAT and AMP differ in many ways. A conservative interpretation of this finding is that regional implicit bias operationalized in terms of the IAT versus AMP reflect distinct constructs. Future research should continue to investigate regional implicit bias operationalized across different measures.

5.3 Connecting Levels of Analysis

Regional intergroup bias research will be advanced to the extent that it can be connected to, and integrated with, intergroup bias research at the individual level. Importantly, making inferences about one level of analysis based on findings at a different level of analysis can lead to incorrect conclusions (Selvin, 1958; Simpson, 1951); in contrast, synthesizing findings from multiple levels of analysis conceptually maps onto Sakaluk's (2016) approach of "exploring small, confirming big" to advance cumulative and replicable psychological research. When questions related to intergroup bias have been investigated at both individual and region levels, the degree to which findings converge versus conflict across levels of analysis can contribute to the development of comprehensive theory that spans levels of analysis.

A few examples already exist in the published intergroup bias literature that include both individual and region levels. In two correlational studies, Bell and colleagues (2021) showed that individuals with higher levels of racial bias report less willingness to adopt a Black child, and

that Black children are adopted from foster care at lower rates in states with higher levels of racial bias. O'Shea and colleagues' (2019) investigated the relationship between racial bias and disease threat, presenting correlational evidence that regional racial bias among Project Implicit visitors is higher in states with higher rates of infectious disease, along with experimental evidence that exposing individuals to disease-related information increased their racial biases. Similarly, Jimenez and colleagues (2021) explored prejudicial reactions to the removal of Native American mascots, presenting correlational evidence that regional bias against Native Americans increased following the removal of Native American sports mascots, along with experimental evidence in the context of a hypothetical legal scenario that exposing individuals to information about a Native American mascot removal increased their punitive judgments against a Native American. Not only does this multi-level approach help to avoid the ecological fallacy by testing rather than assuming correspondence between levels, but it also capitalizes on the strengths of both regional data (i.e., ecological validity) and experimental data (i.e., internal validity) to provide strong evidence for claims.

Whereas O'Shea and colleagues' (2019) and Jimenez and colleagues' (2021) work provides converging evidence across levels of analysis, Marini and colleagues (2013) found diverging evidence: nations with more overweight residents have higher average levels of bias in favor of thin versus overweight people, but overweight individuals have lower levels of bias in favor of thin versus overweight people. Based on the perspective we have proposed in this chapter, one interpretation of this pattern of results is that the processes that are relatively more strongly reflected in regional aggregates of bias (e.g., cultural norms) are positively related to the proportion of thin versus overweight people in the region, whereas the processes that are relatively more strongly reflected in individual estimates of bias (e.g., idiosyncratic experiences;

individual differences) are negatively related to the proportion of thin versus overweight people near the individual. Regardless of its' underlying mechanism, this pattern of results underscores the importance of testing rather than assuming correspondence across levels (Selvin, 1958; Simpson, 1951).

Though there is clear value in testing hypotheses at multiple levels of analysis within the same package of studies, findings can also be synthesized across publications that reflect programs of research that were conducted at different levels. For example, the relationship between racial bias and shooting decisions identified in an experimental laboratory task (e.g., Correll et al., 2002) dovetails with our findings based on police use of force in American communities (Hehman et al., 2018) – which suggests that the relationship between racial bias and lethal force persists across individual and regional levels of analysis (see also Payne & Correll, 2020). That said, findings at one level of analysis can inform findings at another level of analysis only if the construct is valid at both levels of analysis. Measures of explicit and implicit intergroup bias were initially developed based on individual-level theory and validated on individual-level data. However, to date, construct validity evidence only exists for Black-White racial bias as a regional construct (Hehman et al., 2019). Of the 32 published empirical findings we review in this chapter, 22 (68.75%) focus on racial bias, and the other 10 findings are spread across a half dozen intergroup domains. Consequently, additional validation research is necessary before there is sufficient evidence to draw strong conclusions about the construct validity of regional intergroup bias in domains other than race.

As a final point on the topic of relationships between levels of analysis, an open question remains whether measures at one level should predict outcomes at another level, e.g., regional bias predicting individual behavior (or vice versa). A relatively intuitive prediction might be that,

given that levels of analysis are conceptually distinct, conceptual correspondence suggests that regional biases should predict regional outcomes better than individual outcomes, and that individual biases should predict individual outcomes better than regional outcomes. However, as we argue in this chapter, regional bias estimates cancel out measurement error, minimize individual differences, and amplify shared cultural influences. From this perspective, regional bias estimates should be expected to predict individual outcomes if those outcomes reflect a significant degree of shared cultural influences. For example, to the extent that there are cultural differences in food preferences, such that the average American likes their food spicier than does the average German, regional attitudes towards food spiciness should be expected to predict at least some variance in an individual's spicy-food consumption. Of course, this example is not intended to suggest that individual differences in preferences do not influence individual behaviors: some Germans like their food spicier than some Americans do. Instead, to the extent that an individual behavior reflects any substantial degree of cultural influence, then a regional estimate should predict that behavior.

Our own recent research provides an example of individual/regional cross-level effects. In Esposito and Calanchini (2021), we examined the extent to which regional gay friendliness and individuals' sexuality attitudes and sexual orientation are related to migration. We found that straight people migrated as a function of both their individual sexuality attitudes and regional gay friendliness, such that straight people with relatively pro-straight attitudes migrated to places that were low in gay friendliness but straight people with relatively pro-gay attitudes migrated to places that were high in gay friendliness. In contrast, lesbian, gay, and bisexual people migrated as a function of regional gay friendliness, to destinations that were relatively high in gay friendliness, regardless of their own sexuality attitudes. This example demonstrates two possible

cross-level effects: individual and regional factors jointly predicted straight people's migration, whereas regional factors predicted lesbian, gay, and bisexual people's migration. Building on this illustrative demonstration, future research should continue to investigate relationships between biases and outcomes at multiple levels of analysis.

5.4 With Great Statistical Power Comes Great Statistical Responsibility

The datasets used in most regional intergroup bias research to date have been orders of magnitude larger than those typically used in laboratory-based research. Indeed, even the smallest of the publicly-available Project Implicit datasets we report here (i.e., religion) reflects the responses of over 100,000 participants. The primary benefit of these large samples is that, when aggregated to region levels, they can provide very precise estimates of regional bias – which, in turn, increases statistical power to detect effects by reducing measurement error. That said, the large samples often used in regional analyses do not automatically bias results towards statistical significance. Indeed, in the discriminant validity analyses reported by Hehman et al. (2019), Payne et al. (2019), and Riddle and Sinclair (2019), regional bias predicts theoretically-related outcomes but not theoretically-unrelated outcomes – which illustrates that well-powered regional analyses are not necessarily biased towards statistical significance.

Moreover, statistical power is conceptualized differently when the units of analysis reflect aggregated observations. In the context of regional aggregation, statistical power depends in part on the amount of variance between versus within regions, i.e., the intraclass correlation coefficient. Researchers can calculate an “effective sample size” (Rao & Scott, 1992) by dividing the number of regional units by the intraclass correlation coefficient.¹³ For example, a regional

¹³ The equation given in the text is an approximation of the design effect formula (Higgins et al., 2019).

analysis of American states would have a maximum of 50 units of analysis. Current scientific best practices, derived from research on individuals, suggests that a sample of 50 participants is too small to reliably detect most psychological effects (Cohen, 1992; da Silva Frost & Ledgerwood, 2020; Schönbrodt & Perugini, 2013). However, if 20% of the overall variance in a regional measure of prejudice reflects between-state variance, then the effective sample size of the regional analysis ($50/.20 = 250$) would provide roughly the same statistical power as a study with 250 unclustered observations.

To the extent that aggregation increases statistical power by reducing measurement error and increasing the effective sample size, regional intergroup bias effects may easily be distinguished from zero (i.e., statistically significant), even when the effects are very small effects in absolute value. A body of literature characterized by significant effects that are not necessarily large may be undesirable to researchers, practitioners, or policy makers who prioritize the impact (i.e., effect size) of scientific findings. One solution to this situation is to identify a priori a smallest effect size of interest (Lakens et al., 2018) and then only interpret effects of that size or larger. A drawback to this approach is that some statistically-small effects may have societally-large effects (Götz et al., 2021), especially when they affect a single person repeatedly, or when they affect many people simultaneously (Greenwald et al., 2015). Rosenthal (1990) provides the classic illustration of this point: The effect size of aspirin on reducing heart attacks is $r = .035$, which would be considered negligible by conventional effect size standards (e.g., Cohen, 1992) and, thus, smaller than the smallest effect size of interest that most researchers would choose to investigate or interpret. However, regular doses of aspirin would prevent hundreds of thousands of heart attacks (Greenwald et al., 2015), and are overwhelmingly endorsed by physicians. In fact, when doctors originally became aware of this effect, the clinical

trials were halted because the researchers considered it unethical to continue without offering this treatment to participants in the placebo and control groups. Thus, the statistically-small effect of aspirin on reducing heart attacks translates into a societally-large effect.

To the extent that regional intergroup bias reflects shared culture, as we argue, it is especially likely to affect a single person repeatedly, and affect many people simultaneously. Though some effects identified using a regional approach may ultimately turn out to be relatively inconsequential, we argue that it is too early for research in this area to say for certain which effects (and which effect sizes) are meaningful at the society level. Consequently, we do not yet advocate that researchers ignore effects that are below a certain size. Instead, future research into the moderators, mechanisms, and motivations underlying regional intergroup bias may help us to understand when and how small aggregate effects translate into societally-significant outcomes, and will aid the field in determining what is meaningful.

6. Conclusion

As regional intergroup bias research continues to advance, it will contribute to a growing body of research on macropsychological constructs. For example, self-esteem (Bleidorn et al., 2016), ideology (Motyl et al., 2014), and personality (Rentfrow et al., 2008) are all individual-level psychological constructs that have been studied at region levels. In this chapter we have outlined a theoretical perspective specifically focused on regional intergroup bias. Future research should investigate the extent to which specific models, like ours and those of Murphy and colleagues (2013, 2018) and Payne and colleagues (2017a), versus seemingly-more parsimonious, general-purpose models (e.g., Rentfrow et al., 2008) can best account for different aspects of macropsychological constructs.

Research on regional intergroup biases has helped to shed psychological insight on important large-scale outcomes. This small-but-growing area of study has proved especially fruitful for studying phenomena that are infrequent or otherwise difficult to recreate in the laboratory, typically with relatively high degrees of ecological validity, statistical power, and generalizability. Just as science-fiction writer N. K. Jemison set her *Broken Earth* trilogy in a world where “places have a mind of their own” (Jemisin, 2019, p. xiii), we hope that our work will help future regional intergroup bias researchers to connect psychology to places. In highlighting the role of culture in regional intergroup bias, we join a growing body of psychology-focused research recognizing the structural and institutional contributions to intergroup bias. With an eye to the future, we have outlined a new perspective on regional intergroup bias in hopes of inspiring and advancing further research that is theoretically and empirically rigorous and societally consequential.

References

- Allport, G. W. (1954). *The nature of prejudice*. Perseus Books.
- Angrist, J. D., & Pischke, J. S. (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.
- Axt, J. R., Bar-Anan, Y., & Vianello, M. (2020). The relation between evaluation and racial categorization of emotional faces. *Social Psychological and Personality Science, 11*(2), 196-206.
- Balduzzi, S., Rucker, G., & Schwarzer, G. (2019). How to perform a meta-analysis with R: A practical tutorial. *Evidence-Based Mental Health, 22*(4), 153-160.
- Banaji, M. R. (2001). Implicit attitudes can be measured. In H. L. Roediger III, J. S. Nairne, I. E. Neath, & A. M. Surprenant (Eds.), *The nature of remembering: Essays in honor of Robert G. Crowder* (pp. 117-150). American Psychological Association.
- Bar-Anan, Y., & Nosek, B. A. (2014). A comparative investigation of seven indirect attitude measures. *Behavior Research Methods, 46*(3), 668-688. <https://doi.org/10.3758/s13428-013-0410-6>
- Barnes-Holmes, D., Barnes-Holmes, Y., Power, P., Hayden, E., Milne, R., & Stewart, I. (2006). Do you really know what you believe? Developing the implicit Relational Assessment Procedure (IRAP) as a direct measure of implicit beliefs. *The Irish Psychologist, 32*, 169-177.
- Baron, A. S., & Banaji, M. R. (2006). The development of implicit attitudes: Evidence of race evaluations from ages 6 and 10 and adulthood. *Psychological Science, 17*(1), 53-58. <https://doi.org/10.1111/j.1467-9280.2005.01664.x>

- Bell, S. B., Farr, R., Ofosu, E., Hehman, E., & DeWall, C. N. (2021). Implicit bias predicts less willingness and less frequent adoption of Black children more than explicit bias. *The Journal of Social Psychology*, 1-12.
- Berry, D., & Willoughby, M. T. (2017). On the practical interpretability of cross-lagged panel models: Rethinking a developmental workhorse. *Child Development*, 88(4), 1186-1206.
<https://doi.org/10.1111/cdev.12660>
- Bleidorn, W., Schönbrodt, F., Gebauer, J. E., Rentfrow, P. J., Potter, J., & Gosling, S. D. (2016). To live among like-minded others: Exploring the links between person-city personality fit and self-esteem. *Psychological Science*, 27(3), 419-427.
<https://doi.org/10.1177/0956797615627133>
- Brenner, T. (2017). Identification of clusters: An actor-based approach. *Working Papers on Innovation and Space*, 2(17), 1-19.
- Cameron, C. D., Brown-Iannuzzi, J. L., & Payne, B. K. (2012). Sequential priming measures of implicit social cognition: A meta-analysis of associations with behavior and explicit attitudes. *Personality and Social Psychology Review*, 16(4), 330-350.
<https://doi.org/10.1177/1088868312440047>
- Cao, J., & Banaji, M. R. (2016). The base rate principle and the fairness principle in social judgment. *Proceedings of the National Academy of Sciences*, 113(27), 7475-7480.
- Carpenter, T.P., Gödderz, A., & Lai, C. K. Individual differences in implicit bias can be measured reliably by administering the same Implicit Association Test multiple times.
Unpublished manuscript.

- Charlesworth, T. E., & Banaji, M. R. (2019). Patterns of implicit and explicit attitudes: I. Long-term change and stability from 2007 to 2016. *Psychological Science*, *30*(2), 174-192.
<https://doi.org/10.1177/0956797618813087>
- Cheryan, S., Plaut, V. C., Davies, P. G., & Steele, C. M. (2009). Ambient belonging: How stereotypical cues impact gender participation in computer science. *Journal of Personality and Social Psychology*, *97*(6), 1045-1060. <https://doi.org/10.1037/a0016239>
- Chin, M. J., Quinn, D. M., Dhaliwal, T. K., & Lovison, V. S. (2020). Bias in the air: A nationwide exploration of teachers' implicit racial attitudes, aggregate bias, and student outcomes. *Educational Researcher*, *49*(8), 566-578.
<https://doi.org/10.3102%2F0013189X20937240>
- Cohen, J. (1992). A power primer. *Psychological Bulletin*, *112*(1), 155-159.
<https://doi.org/10.1037/0033-2909.112.1.155>
- Connor, P., & Evers, E. R. (2020). The bias of individuals (in crowds): Why implicit bias is probably a noisily measured individual-level construct. *Perspectives on Psychological Science*, *15*(6), 1329-1345. <https://doi.org/10.1177/1745691620931492>
- Conrey, F. R., Sherman, J. W., Gawronski, B., Hugenberg, K., & Groom, C. (2005). Separating multiple processes in implicit social cognition: The quad model of implicit task performance. *Journal of Personality and Social Psychology*, *89*, 469-487.
<https://doi.org/10.1037/0022-3514.89.4.469>
- Cooley, E., & Payne, B. K. (2017). Using groups to measure intergroup prejudice. *Personality and Social Psychology Bulletin*, *43*(1), 46-59. <https://doi.org/10.1177/0146167216675331>

- Correll, J., Park, B., Judd, C. M., & Wittenbrink, B. (2002). The police officer's dilemma: Using ethnicity to disambiguate potentially threatening individuals. *Journal of Personality and Social Psychology, 83*(6), 1314–1329. <https://doi.org/10.1037/0022-3514.83.6.1314>
- Crandall, C. S., & Eshleman, A. (2003). A justification-suppression model of the expression and experience of prejudice. *Psychological Bulletin, 129*(3), 414-446.
- Cunningham, W. A., Preacher, K. J., & Banaji, M. R. (2001). Implicit attitude measures: Consistency, stability, and convergent validity. *Psychological Science, 12*(2), 163-170. <https://doi.org/10.1111/1467-9280.00328>
- Cunningham, W. A., Zelazo, P. D., Packer, D. J., & Van Bavel, J. J. (2007). The iterative reprocessing model: A multilevel framework for attitudes and evaluation. *Social Cognition, 25*(5), 736-760. <https://doi.org/10.1521/soco.2007.25.5.736>
- da Silva Frost A., & Ledgerwood, A. (2020) Calibrate your confidence in research findings: A tutorial on improving research methods and practices. *Journal of Pacific Rim Psychology, 14*, e14.
- Daumeyer, N. M., Rucker, J. M., & Richeson, J. A. (2017). Thinking structurally about implicit bias: Some peril, lots of promise, *Psychological Inquiry, 28*(4), 258-261. <https://doi.org/10.1080/1047840X.2017.1335568>
- De Houwer, J. (2003). The Extrinsic Affective Simon Task. *Experimental Psychology, 50*, 77–85.
- Devine, P. G., Plant, E. A., Amodio, D. M., Harmon-Jones, E., & Vance, S. L. (2002). The regulation of explicit and implicit race bias: The role of motivations to respond without prejudice. *Journal of Personality and Social Psychology, 82*(5), 835-848. <https://doi.org/10.1037//0022-3514.82.5.835>

- Devos, T., & Sadler, M. (2019). Context diversity predicts the extent to which the American identity is implicitly associated with Asian Americans and European Americans. *Asian American Journal of Psychology, 10*(2), 182–193. <https://doi.org/10.1037/aap0000149>
- Devos, T., Sadler, M., Perry, D., & Yogeeswaran, K. (2021). Temporal fluctuations in context ethnic diversity over three decades predict implicit national inclusion of Asian Americans. *Group Processes & Intergroup Relations, 24*(1), 3-25. <https://doi.org/10.1177/1368430219887440>
- Ebert, T., Götz, F. M., Gladstone, J. J., Müller, S. R., & Matz, S. C. (2020). Spending reflects not only who we are but also who we are around: The joint effects of individual and geographic personality on consumption. *Journal of Personality and Social Psychology*. <https://doi.org/10.1037/pspp0000344>
- Ekstrom, P. D., Le Forestier, J. M., & Lai, C. K. (2021). Racial demographics explain the link between racial disparities in traffic stops and county-level racial attitudes. *Psychological Science*. <https://doi.org/10.31234/osf.io/znhuf>
- Esposito, E., & Calanchini, J. (2021). Gay Migration Versus Person-Environment Fit: A Regional Analysis. Unpublished manuscript.
- Fazio, R. H., Chen, J. M., McDonel, E. C., & Sherman, S. J. (1982). Attitude accessibility, attitude-behavior consistency, and the strength of the object-evaluation association. *Journal of Experimental Social Psychology, 18*(4), 339-357. [https://doi.org/10.1016/0022-1031\(82\)90058-0](https://doi.org/10.1016/0022-1031(82)90058-0)
- Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline?. *Journal of*

- Personality and Social Psychology*, 69(6), 1013-1027. <https://doi.org/10.1037/0022-3514.69.6.1013>
- Fazio, R. H., & Towles-Schwen, T. (1999). The MODE model of attitude-behavior processes. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 97–116). The Guilford Press.
- Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin*, 132(5), 692-731. <https://doi.org/10.1037/0033-2909.132.5.692>
- Forscher, P. S., Lai, C. K., Axt, J. R., Ebersole, C. R., Herman, M., Devine, P. G., & Nosek, B. A. (2019). A meta-analysis of procedures to change implicit measures. *Journal of Personality and Social Psychology*, 117(3), 522-559. <https://doi.org/10.1037/pspa0000160>
- Garcia-Marques, L., Santos, A. S., Mackie, D. M., Hagá, S., & Palma, T. A. (2017) Cognitive malleability and the wisdom of independent aggregation. *Psychological Inquiry*, 28(4), 262-267. <https://doi.org/10.1080/1047840X.2017.1373558>
- Gawronski, B. & Bodenhausen, G. V. (2017) Beyond persons and situations: An interactionist approach to understanding implicit bias. *Psychological Inquiry*, 28(4), 268-272. <https://doi.org/10.1080/1047840X.2017.1373546>
- Gawronski, B., Morrison, M., Phills, C. E., & Galdi, S. (2017). Temporal stability of implicit and explicit measures: A longitudinal analysis. *Personality and Social Psychology Bulletin*, 43(3), 300-312. <https://doi.org/10.1177/0146167216684131>
- Gawronski, B., & De Houwer, J. (2014). Implicit measures in social and personality psychology. In H. T. Reis, & C. M. Judd (Eds.), *Handbook of research methods in social and*

- personality psychology* (2nd edition, pp. 283-310). New York, NY: Cambridge University Press.
- Götz, F. M., Ebert, T., Gosling, S. D., Obschonka, M., Potter, J., & Rentfrow, P. J. (2021). Local housing market dynamics predict rapid shifts in cultural openness: A 9-year study across 199 cities. *American Psychologist*, *76*(6), 947–961.
- Götz, F. M., Gosling, S. D., & Rentfrow, P. J., (2021). Small effects: The indispensable foundation for a cumulative psychological science. *Perspectives on Psychological Science*. <https://doi.org/10.1177/1745691620984483>
- Granger, C. W. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, *37*(3), 424-438. <https://doi.org/10.2307/1912791>
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, *102*(1), 4–27. <https://doi.org/10.1037/0033-295x.102.1.4>
- Greenwald, A. G., Banaji, M. R., & Nosek, B. A. (2015). Statistically small effects of the Implicit Association Test can have societally large effects. *Journal of Personality and Social Psychology*, *108*(4), 553–561. <https://doi.org/10.1037/pspa0000016>
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology*, *74*(6), 1464-1480. <https://doi.org/10.1037//0022-3514.74.6.1464>
- Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the Implicit Association Test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology*, *85*(2), 197–216. <https://doi.org/10.1037/0022-3514.85.2.197>

- Greenwald, A. G., Poehlman, T. A., Uhlmann, E. L., & Banaji, M. R. (2009). Understanding and using the Implicit Association Test: III. Meta-analysis of predictive validity. *Journal of Personality and Social Psychology*, *97*(1), 17-41. <https://doi.org/10.1037/a0015575>
- Grimm, K. J., Helm, J., Rodgers, D., & O'Rourke, H. (2021). Analyzing cross-lag effects: A comparison of different cross-lag modeling approaches. *New Directions for Child and Adolescent Development*, *2021*(175), 11-33.
- Hamilton, D. L. (1981). Stereotyping and intergroup behavior: Some thoughts on the cognitive approach. In D. L. Hamilton (Ed.), *Cognitive processes in stereotyping and intergroup behavior* (pp. 333–353). Erlbaum.
- Hannay, J. W. & Payne, B. K. (2021). Implicit bias in character and context: The role of aggregation. Unpublished manuscript.
- Hatzenbuehler, M.L., Weissman, D.G., McKetta, S., Lattanner, M.R., Ford, J.V., Barch, D.M., & McLaughlin, K.A. (2021). Smaller hippocampal volume among Black and Latinx youth living in high-stigma contexts. *Journal of the American Academy of Child and Adolescent Psychiatry*. <https://doi.org/10.1016/j.jaac.2021.08.017>
- Hehman, E., Calanchini, J., Flake, J. K., & Leitner, J. B. (2019). Establishing construct validity evidence for region-level measures of explicit and implicit racial bias. *Journal of Experimental Psychology: General*, *148*(6), 1022-1040. <https://doi.org/10.1037/xge0000623>
- Hehman, E., Flake, J. K., & Calanchini, J. (2018). Disproportionate use of lethal force in policing is associated with regional racial biases of residents. *Social Psychological and Personality Science*, *9*(4), 393-401. <https://doi.org/10.1177/1948550617711229>

- Hehman, E., Oforu, E. K., & Calanchini, J. (2020). Using environmental features to maximize prediction of regional intergroup bias. *Social Psychological and Personality Science*, *12*(2), 156-164. <https://doi.org/10.1177/1948550620909775>
- Higgins, J. P., Thomas, J., Chandler, J., Cumpston, M., Li, T., Page, M. J., & Welch, V. A. (Eds.). (2019). *Cochrane handbook for systematic reviews of interventions*. John Wiley & Sons.
- Hofmann, W., Gawronski, B., Gschwendner, T., Le, H., & Schmitt, M. (2005). A meta-analysis on the correlation between the Implicit Association Test and explicit self-report measures. *Personality and Social Psychology Bulletin*, *31*(10), 1369-1385. <https://doi.org/10.1177/0146167205275613>
- Hoover, J., & Dehghani, M. (2019). The big, the bad, and the ugly: Geographic estimation with flawed psychological data. *Psychological Methods*, *25*(4), 412-429. <https://doi.org/10.1037/met0000240>
- Hornsey, M. J. (2008). Social identity theory and self-categorization theory: A historical review. *Social and Personality Psychology Compass*, *2*(1), 204-222. <https://doi.org/10.1111/j.1751-9004.2007.00066.x>
- Jemisin, N. K. (2019). *How long 'til black future month?* Orbit Books.
- Johnson, B. T., Cromley, E. K., & Natasza Marrouch (2017). Spatiotemporal meta-analysis: Reviewing health psychology phenomena over space and time. *Health Psychology Review*, *11*(3), 280-291. <https://doi.org/10.1080/17437199.2017.1343679>
- Johnson, D. J., & Chopik, W. J. (2019). Geographic variation in the black-violence stereotype. *Social Psychological and Personality Science*, *10*(3), 287-294. <https://doi.org/10.1177/1948550617753522>

- Jost, J. T., Banaji, M. R., & Nosek, B. A. (2004). A decade of system justification theory: Accumulated evidence of conscious and unconscious bolstering of the status quo. *Political Psychology, 25*(6), 881-919. <https://doi.org/10.1111/j.1467-9221.2004.00402.x>
- Kane, M. T. (2013). Validating the interpretations and uses of test scores. *Journal of Educational Measurement, 50*(1), 1–73. <https://doi.org/10.1111/jedm.12000>
- Karpinski, A., & Hilton, J. L. (2001). Attitudes and the Implicit Association Test. *Journal of Personality and Social Psychology, 81*(5), 774–788. <https://doi.org/10.1037/0022-3514.81.5.774>
- Kondo, K. (2016). Hot and cold spot analysis using Stata. *The Stata Journal, 16*(3), 613-631. <https://doi.org/10.1177/1536867X1601600304>
- Kunst, J. R., Fischer, R., Sidanius, J., & Thomsen, L. (2017). Preferences for group dominance track and mediate the effects of macro-level social inequality and violence across societies. *Proceedings of the National Academy of Sciences, 114*(21), 5407-5412. <https://doi.org/10.1073/pnas.1616572114>
- Kurdi, B., & Banaji, M. R. (2017) Reports of the death of the individual difference approach to implicit social cognition may be greatly exaggerated: A commentary on Payne, Vuletich, and Lundberg. *Psychological Inquiry, 28*(4), 281-287. <https://doi.org/10.1080/1047840X.2017.1373555>
- Kurdi, B., Seitchik, A. E., Axt, J. R., Carroll, T. J., Karapetyan, A., Kaushik, N., Tomczsko, D., Greenwald, A. G., & Banaji, M. R. (2019). Relationship between the Implicit Association Test and intergroup behavior: A meta-analysis. *American Psychologist, 74*(5), 569-586. <https://doi.org/10.1037/amp0000364>

- Kwan, M. (2012). The uncertain geographic context problem. *Annals of The Association of American Geographers, 102*(5), 958-968. <https://doi.org/10.1080/00045608.2012.687349>
- Lai, C. K., Skinner, A. L., Cooley, E., Murrar, S., Brauer, M., Devos, T., Calanchini, J., Xiao, Y., J., Pedram, C., Marshburn, C. K., Simon, S., Blanchar, J. C., Joy-Gaba, J. A., Conway, J., Redford, L., Klein, R. A., Roussos, G., Schellhaas, F. M. H., Burns, M., ... Nosek, B. A. (2016). Reducing implicit racial preferences: II. Intervention effectiveness across time. *Journal of Experimental Psychology: General, 145*(8), 1001-1016. <https://doi.org/10.1037/xge0000179>
- Lakens, D., Scheel, A. M., & Isager, P. M. (2018). Equivalence testing for psychological research: A tutorial. *Advances in Methods and Practices in Psychological Science, 1*(2), 259-269. <https://doi.org/10.1177/2515245918770963>
- Latané, B. (1981). The psychology of social impact. *American Psychologist, 36*(4), 343-356. <https://doi.org/10.1037/0003-066X.36.4.343>
- Lavery, M. R., Acharya, P., Sivo, S. A., & Xu, L. (2019). Number of predictors and multicollinearity: What are their effects on error and bias in regression?. *Communications in Statistics-Simulation and Computation, 48*(1), 27-38. <https://doi.org/10.1080/03610918.2017.1371750>
- 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. <https://doi.org/10.1177/0956797616658450>

- Leitner, J. B., Hehman, E., Ayduk, O., & Mendoza-Denton, R. (2016b). Racial bias is associated with ingroup death rate for Blacks and Whites: Insights from Project Implicit. *Social Science & Medicine*, *170*, 220-227. <https://doi.org/10.1016/j.socscimed.2016.10.007>
- Leitner, J. B., Hehman, E., & Snowden, L. R. (2018). States higher in racial bias spend less on disabled Medicaid enrollees. *Social Science & Medicine*, *208*, 150-157. <https://doi.org/10.1016/j.socscimed.2018.01.013>
- Levy, B. R., & Banaji, M. R. (2002). Implicit ageism. In T. D. Nelson (Ed.), *Ageism: Stereotyping and prejudice against older persons* (pp. 49-75). The MIT Press.
- Lewin, K. (1951). Problems of research in social psychology. In D. Cartwright (Ed.), *Field theory in social science: Selected theoretical papers* (pp. 155-169). Harper & Row.
- Maas, C. J., & Hox, J. J. (2005). Sufficient sample sizes for multilevel modeling. *Methodology*, *1*(3), 86-92.
- Machery, E. (2017). Do indirect measures of biases measure traits or situations?. *Psychological Inquiry*, *28*, 288-291. <https://doi.org/10.1080/1047840X.2017.1373557>
- Marini, M., Sriram, N., Schnabel, K., Maliszewski, N., Devos, T., Ekehammar, B., Wiers, R., HuaJian, C., Somogyi, M., Shiomura, K., Schnall, S., Neto, F., Bar-Anan, Y., Vianello, M., Ayala, A., Dorantes, G., Park, J., Kesebir, S., Pereira, A., ... Nosek, B. A. (2013). Overweight people have low levels of implicit weight bias, but overweight nations have high levels of implicit weight bias. *PLoS ONE*, *8*(12), e83543. <https://doi.org/10.1371/journal.pone.0083543>
- Meissner, F., & Rothermund, K. (2013). Estimating the contributions of associations and recoding in the Implicit Association Test: The ReAL model for the IAT. *Journal of Personality and Social Psychology*, *104*(1), 45-69. <https://doi.org/10.1037/a0030734>

- Miller, C. T., Varni, S. E., Solomon, S. E., DeSarno, M. J., & Bunn, J. Y. (2016). Macro-level implicit HIV prejudice and the health of community residents with HIV. *Health Psychology, 35*, 807-815. <https://doi.org/10.1037/hea0000314>
- Moors, A., & De Houwer, J. (2006). Automaticity: a theoretical and conceptual analysis. *Psychological Bulletin, 132*(2), 297.
- Moran, P. A. (1950). Notes on continuous stochastic phenomena. *Biometrika, 37*(1/2), 17-23. <https://doi.org/10.2307/2332142>
- Moshontz, H., Campbell, L., Ebersole, C. R., IJzerman, H., Urry, H. L., Forscher, P. S., Grahe, J. E., McCarthy, R. J., Musser, E. D., Antfolk, J., Castille, C. M., Evans, T. R., Fielder, S., Flake, J. K., Forero, D. A., Janssen, S. M. J., Keene, J. R., Protzko, J., Aczel, B., ... Chartier, C. R. (2018). The Psychological Science Accelerator: Advancing psychology through a distributed collaborative network. *Advances in Methods and Practices in Psychological Science, 1*(4), 501-515. <https://doi.org/10.1177/2515245918797607>
- Motyl, M., Iyer, R., Oishi, S., Trawalter, S., & Nosek, B. A. (2014). How ideological migration geographically segregates groups. *Journal of Experimental Social Psychology, 51*, 1-14. <https://doi.org/10.1016/j.jesp.2013.10.010>
- Murphy, M. C., Kroeper, K. M., & Ozier, E. M. (2018). Prejudiced places: How contexts shape inequality and how policy can change them. *Policy Insights from the Behavioral and Brain Sciences, 5*(1), 66-74. <https://doi.org/10.1177/2372732217748671>
- Murphy, M. C., Steele, C. M., & Gross, J. J. (2007). Signaling threat: How situational cues affect women in math, science, and engineering settings. *Psychological Science, 18*(10), 879-885. <https://doi.org/10.1111/j.1467-9280.2007.01995.x>

- Murphy, M. C., & Walton, G. M. (2013). From prejudiced people to prejudiced places: A social-contextual approach to prejudice. In C. Stangor & C. S. Crandall (Eds.), *Stereotyping and Prejudice* (pp. 181-204). Psychology Press.
- Nadarevic, L., & Erdfelder, E. (2011). Cognitive processes in implicit attitude tasks: An experimental validation of the Trip Model. *European Journal of Social Psychology*, *41*(2), 254-268. <https://doi.org/10.1002/ejsp.776>
- Nosek B. A. (2005). Moderators of the relationship between implicit and explicit evaluation. *Journal of Experimental Psychology. General*, *134*(4), 565–584. <https://doi.org/10.1037/0096-3445.134.4.565>
- Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2005). Understanding and using the Implicit Association Test: II. Method variables and construct validity. *Journal of Personality and Social Psychology*, *31*(2), 166 –180. <https://doi.org/10.1177/0146167204271418>
- Nosek, B. A., Smyth, F. L., Hansen, J. J., Devos, T., Lindner, N. M., Ranganath, K. A., Tucker Smith, C., Olson, K. R., Chugh, D., Greenwald, A. G., & Banaji, M. R. (2007). Pervasiveness and correlates of implicit attitudes and stereotypes. *European Review of Social Psychology*, *18*(1), 36-88. <https://doi.org/10.1080/10463280701489053>
- Nosek, B. A., Smyth, F. L., Sriram, N., Lindner, N. M., Devos, T., Ayala, A., Bar-Anan, Y., Bergh, R., Cai, H., Gonsalkorale, K., Kesebir, S., Maliszewski, N., Neto, F., Olli, E., Park, J., Schnabel, K., Shiomura, K., Tulbure, B. T., Wiers, R. W., ... Greenwald, A. G. (2009). National differences in gender–science stereotypes predict national sex differences in science and math achievement. *Proceedings of the National Academy of Sciences*, *106*(26), 10593-10597. <https://doi.org/10.1073/pnas.0809921106>
- Nunnally, J. (1970). *Introduction to psychological measurement*. McGraw-Hill.

- Ofori, E. K., Chambers, M. K., Chen, J. M., & Hehman, E. (2019). Same-sex marriage legalization associated with reduced implicit and explicit antigay bias. *Proceedings of the National Academy of Sciences, 116*(18), 8846-8851.
<https://doi.org/10.1073/pnas.1806000116>
- Oishi, S., & Graham, J. (2010). Social ecology: Lost and found in psychological science. *Perspectives on Psychological Science, 5*(4), 356-377.
<https://doi.org/10.1177/1745691610374588>
- Ojala, M., & Garriga, G. C. (2010). Permutation tests for studying classifier performance. *Journal of Machine Learning Research, 11*(6), 1833-1863.
- Olson, M. A., & Fazio, R. H. (2003). Relations between implicit measures of prejudice: What are we measuring? *Psychological Science, 14*, 636-639. https://doi.org/10.1046/j.0956-7976.2003.psci_1477.x
- Olson, M. A., & Fazio, R. H. (2004). Reducing the influence of extrapersonal associations on the Implicit Association Test: Personalizing the IAT. *Journal of Personality and Social Psychology, 86*(5), 653-667. <https://doi.org/10.1037/0022-3514.86.5.653>
- Orchard, J., & Price, J. (2017). County-level racial prejudice and the black-white gap in infant health outcomes. *Social Science & Medicine, 181*, 191-198.
<https://doi.org/10.1016/j.socscimed.2017.03.036>
- O'Shea, B. A., Watson, D. G., Brown, G. D., & Fincher, C. L. (2019). Infectious disease prevalence, not race exposure, predicts both implicit and explicit racial prejudice across the United States. *Social Psychological and Personality Science, 11*(3), 345-355.
<https://doi.org/10.1177/1948550619862319>

- Oswald, F. L., Mitchell, G., Blanton, H., Jaccard, J., & Tetlock, P. E. (2013). Predicting ethnic and racial discrimination: A meta-analysis of IAT criterion studies. *Journal of Personality and Social Psychology, 105*(2), 171-192. <https://doi.org/10.1037/a0032734>
- Payne, B. K., Burkley, M. A., & Stokes, M. B. (2008). Why do implicit and explicit attitude tests diverge? The role of structural fit. *Journal of Personality and Social Psychology, 94*(1), 16–31. <https://doi.org/10.1037/0022-3514.94.1.16>
- Payne, B. K., & Correll, J. (2020). Race, weapons, and the perception of threat. *Advances in Experimental Social Psychology, 62*, 1-50. <https://doi.org/10.1016/bs.aesp.2020.04.001>
- Payne, B. K., Cheng, C. M., Govorun, O., & Stewart, B. D. (2005). An inkblot for attitudes: Affect misattribution as implicit measurement. *Journal of Personality and Social Psychology, 89*(3), 277-293. <https://doi.org/10.1037/0022-3514.89.3.277>
- Payne, B. K., & Hannay, J. W. (2021). Implicit bias reflects systemic racism. *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2021.08.001>
- Payne, B. K., Vuletich, H. A., & Brown-Iannuzzi, J. L. (2019). Historical roots of implicit bias in slavery. *Proceedings of the National Academy of Sciences, 116*(24), 11693-11698. <https://doi.org/10.1073/pnas.1818816116>
- Payne, B. K., Vuletich, H. A., & Lundberg, K. B. (2017a). The bias of crowds: How implicit bias bridges personal and systemic prejudice. *Psychological Inquiry, 28*(4), 233-248. <https://doi.org/10.1080/1047840X.2017.1335568>
- Payne, B. K., Vuletich, H. A., & Lundberg, K. B. (2017b) Flipping the script on implicit bias research with the bias of crowds. *Psychological Inquiry, 28*(4), 306-311. <https://doi.org/10.1080/1047840X.2017.1380460>

- Pearman, F. A. (2021). Collective racial bias and the Black-White test score gap. *Race and Social Problems*, 1-10. <https://doi.org/10.26300/58zh-8v92>
- Pettigrew, T. F. (1997). Personality and social structure: Social psychological contributions. In R. Hogan, J. A. Johnson, & S. R. Briggs (Eds.), *Handbook of personality psychology* (pp. 417- 438). Academic Press.
- Pincus, F. L. (1996). Discrimination comes in many forms: Individual, institutional, and structural. *American Behavioral Scientist*, 40(2), 186-194.
<https://doi.org/10.1177/0002764296040002009>
- Plant, E. A., & Devine, P. G. (1998). Internal and external motivation to respond without prejudice. *Journal of Personality and Social Psychology*, 75(3), 811–832.
<https://doi.org/10.1037/0022-3514.75.3.811>
- Price, M., McKetta, S., Weisz, J., Ford, J., Lattanner, M., Skov, H., Wolock, E., & Hatzenbuehler, M. (2021b). Cultural sexism moderates efficacy of psychological therapy for girls: Results from a spatial meta-analysis. *Clinical Psychology: Science and Practice*.
- Price, M. A., Weisz, J. R., McKetta, S., Hollinsaid, N. L., Lattanner, M. R., Reid, A. E., & Hatzenbuehler, M. L. (2021a). Meta-analysis: Are psychotherapies less effective for black youth in communities with higher levels of anti-black racism?. *Journal of the American Academy of Child & Adolescent Psychiatry*.
<https://doi.org/10.1016/j.jaac.2021.07.808>
- Rae, J. R., & Greenwald, A. G. (2017) Persons or situations? Individual differences explain variance in aggregated implicit race attitudes. *Psychological Inquiry*, 28(4), 297-300.
<https://doi.org/10.1080/1047840X.2017.1373548>

- Rae, J. R., Newheiser, A. K., & Olson, K. R. (2015). Exposure to racial out-groups and implicit race bias in the United States. *Social Psychological and Personality Science*, 6(5), 535-543. <https://doi.org/10.1177/1948550614567357>
- Ramakrishnan, K., Lee, J., Lee, T., & Wong, J. (2017). *National Asian American Survey (NAAS) 2016 Pre-Election Survey*. National Asian American Survey. <http://naasurvey.com/data/>
- Rao, J. N. K., & Scott, A. J. (1992). A simple method for the analysis of clustered binary data. *Biometrics*, 577-585.
- Reid, A. E., Dovidio, J. F., Ballester, E., & Johnson, B. T. (2014). HIV prevention interventions to reduce sexual risk for African Americans: The influence of community-level stigma and psychological processes. *Social Science & Medicine*, 103, 118-125.
- Rentfrow, P. J. (2010). Statewide differences in personality: Toward a psychological geography of the United States. *American Psychologist*, 65(6), 548-558.
<https://doi.org/10.1037/a0018194>
- Rentfrow, P. J., Gosling, S. D., & Potter, J. (2008). A theory of the emergence, persistence, and expression of geographic variation in psychological characteristics. *Perspectives on Psychological Science*, 3(5), 339-369. <https://doi.org/10.1111/j.1745-6924.2008.00084.x>
- Rentfrow, P. J., & Jokela, M. (2016). Geographical psychology: The spatial organization of psychological phenomena. *Current Directions in Psychological Science*, 25(6), 393-398.
<https://doi.org/10.1177/0963721416658446>
- Riddle, T., & Sinclair, S. (2019). Racial disparities in school-based disciplinary actions are associated with county-level rates of racial bias. *Proceedings of the National Academy of Sciences*, 116(17), 8255-8260.

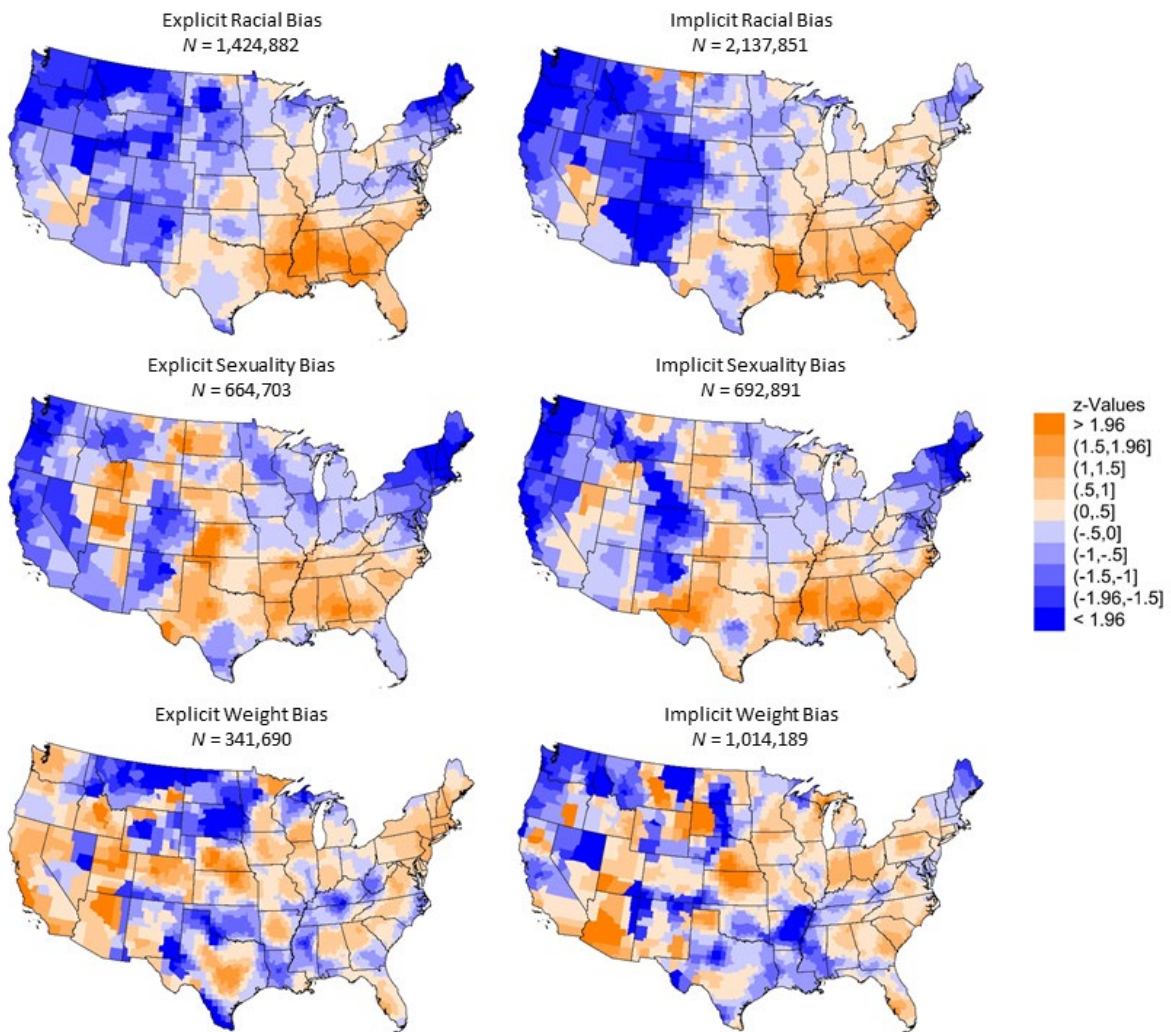
- Robinson, W. S. (1950). Ecological correlations and the behavior of individuals. *American Sociological Review*, *15*(3), 351–357. <https://doi.org/10.2307/2087176>
- Roese, N. J., & Jamieson, D. W. (1993). Twenty years of bogus pipeline research: A critical review and meta-analysis. *Psychological Bulletin*, *114*(2), 363–375. <https://doi.org/10.1037/0033-2909.114.2.363>
- Rosenthal, R. (1990). How are we doing in soft psychology? *American Psychologist*, *45*(6), 775-777. <https://doi.org/10.1037/0003-066X.45.6.775>
- Rushton, J. P., Brainerd, C. J., & Pressley, M. (1983). Behavioral development and construct validity: The principle of aggregation. *Psychological Bulletin*, *94*(1), 18-38. <https://doi.org/10.1037/0033-2909.94.1.18>
- Sakaluk, J. K. (2016). Exploring small, confirming big: An alternative system to the new statistics for advancing cumulative and replicable psychological research. *Journal of Experimental Social Psychology*, *66*, 47-54.
- Samayoa, J. A. G., & Fazio, R. H. (2017) Who starts the wave? Let's not forget the role of the individual. *Psychological Inquiry*, *28*(4), 273-277. <https://doi.org/10.1080/1047840X.2017.1373554>
- Schönbrodt, F. D., & Perugini, M. (2013). At what sample size do correlations stabilize?. *Journal of Research in Personality*, *47*(5), 609-612.
- Sears, D. O. (1986). College sophomores in the laboratory: Influences of a narrow data base on social psychology's view of human nature. *Journal of Personality and Social Psychology*, *51*(3), 515-530. <https://doi.org/10.1037/0022-3514.51.3.515>
- Selvin, H. C. (1958). Durkheim's suicide and problems of empirical research. *American Journal of Sociology*, *63*(6), 607-619. <https://doi.org/10.1086/222356>

- Simpson, Edward H. (1951). The Interpretation of Interaction in Contingency Tables. *Journal of the Royal Statistical Society, Series B*, 13, 238–241.
- Smaldino, P. E. (2014). Group-level traits emerge. *Behavioral and Brain Sciences*, 37(3), 281-295. <https://doi.org/10.1017/S0140525X13003531>
- Spearman, C. (1904). The proof and measurement of association between two things. *The American Journal of Psychology*, 15, 72–101. <https://doi.org/10.1037/11491-005>
- Stahl, C., & Degner, J. (2007). Assessing automatic activation of valence: A multinomial model of EAST performance. *Experimental Psychology*, 54(2), 99-112. 10.1027/1618-3169.54.2.99
- Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review*, 8(3), 220-247. https://doi.org/10.1207/s15327957pspr0803_1
- Stelter, M., Essien, I., Sander, C., & Degner, J. (2021). Racial bias in police traffic stops: White residents' county-level prejudice and stereotypes are related to disproportionate stopping of Black drivers. *Psychological Science*.
- Triandis, H. C. (2002). Subjective culture. *Online Readings in Psychology and Culture*, 2(2), 1-12. <http://dx.doi.org/10.9707/2307-0919.1021>
- Usami, S., Murayama, K., & Hamaker, E. L. (2019). A unified framework of longitudinal models to examine reciprocal relations. *Psychological Methods*, 24(5), 637.
- Varnum, M. E., & Grossmann, I. (2017). Cultural change: The how and the why. *Perspectives on Psychological Science*, 12(6), 956-972. <https://doi.org/10.1177/1745691617699971>
- Vuletich, H. A., & Payne, B. K. (2019). Stability and change in implicit bias. *Psychological Science*, 30(6), 854-862. <https://doi.org/10.1177/0956797619844270>

- Walton, G. M., & Cohen, G. L. (2007). A question of belonging: Race, social fit, and achievement. *Journal of Personality and Social Psychology, 92*(1), 82-96.
<https://doi.org/10.1037/0022-3514.92.1.82>
- Weisbuch, M., Pauker, K., & Ambady, N. (2009). The subtle transmission of race bias via televised nonverbal behavior. *Science, 326*(5960), 1711-1714.
- Wightman, L. F. (1998). An examination of sex differences in LSAT scores from the perspective of social consequences. *Applied Measurement in Education, 11*(3), 255-277.
- Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review, 107*(1), 101-126. <https://doi.org/10.1037/0033-295X.107.1.101>
- Zerhouni, O., Rougier, M., & Muller, D. (2016). “Who (really) is Charlie?” French cities with lower implicit prejudice toward Arabs demonstrated larger participation rates in Charlie Hebdo rallies. *International Review of Social Psychology, 29*(1), 69-76.
<http://dx.doi.org/10.5334/irsp.50>
- Zestcott, C. A., Ruiz, J. M., Tietje, K. R., & Stone, J. (2021). The Relationship Between Racial Prejudice and Cardiovascular Disease Mortality Risk at the State and County Level. *Annals of Behavioral Medicine.*

Figure 1

Maps of explicit (left) and implicit (right) regional intergroup biases, as reflected in the responses of Project Implicit visitors





Note. Maps generated using distance-based weighting (Brenner, 2017; Ebert et al., 2021). Specifically, we used the most fine-grained geographical information available in the data (i.e., visitor's county of residence) and calculated a score for each geographical unit that is based on all observations in the data. Thus, each observation is weighted according to its distance to the target geographical unit. To depict distributional patterns across the U.S., we defined a log-logistic weighting scheme in which observations that are up to 20 miles away from the target geographical unit receive a weight of nearly one, observations that are 50 miles away receive a weight of 0.5 and observations further away than 80 miles receive a weight of nearly zero.

Figure 2

Meta-analyses of the correlation between explicit and implicit measures of state-level intergroup bias

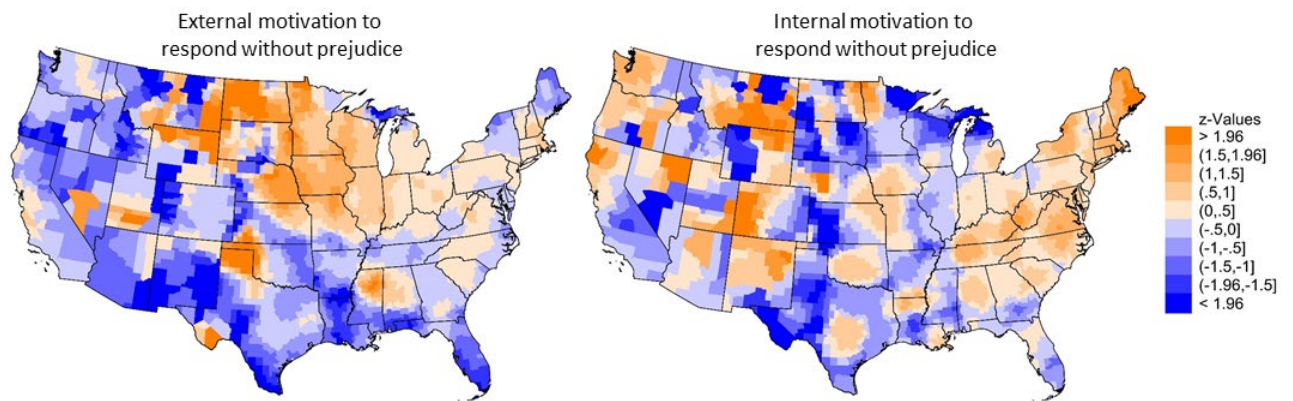
2a.



Note. Top panel (2a) reflects attitude domains, and bottom panel (2b) reflects stereotype domains. Squares reflect point estimates of state-level correlations between explicit and implicit bias, and error bars reflect 95% confidence intervals.

Figure 3

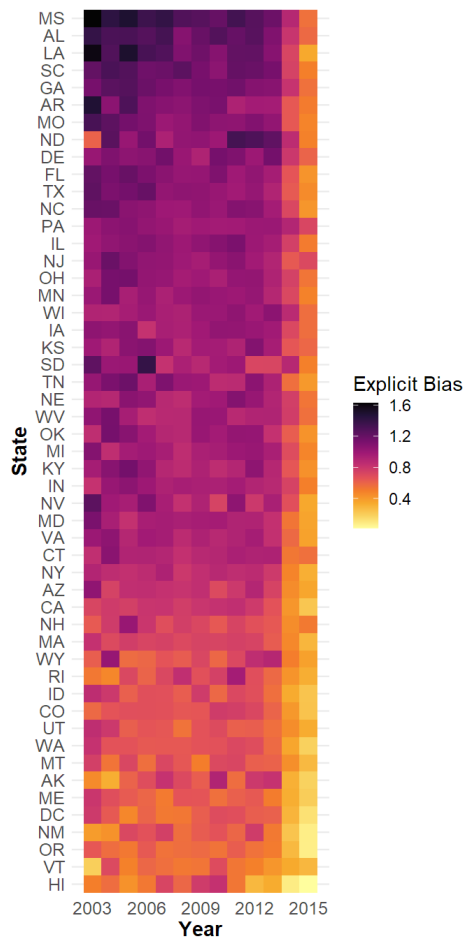
Maps reflecting regional variation in external (left) and internal (right) motivations to respond without prejudice



Note. Maps of the motivations to respond without prejudice (Plant & Devine, 1998), based on the responses of $N = 100,262$ Project Implicit visitors. Maps were generated using the same distance-based weighting approach as in Figure 1 (Brenner, 2017; Ebert et al., 2021).

Figure 4


Explicit preferences for heterosexual versus gay people over time, by U.S. state



Note. This figure illustrates the substantial between- and within-state heterogeneity in explicit sexuality bias over time. Higher values reflect greater evaluative preference for heterosexual versus gay people.

Figure 5

Regional variation in mental health providers in the United States as a function of sexuality bias



Note. Regional units are core-based statistical areas. Darker blue regions reflect fewer mental health providers. Darker orange regions reflect more anti-gay sexuality bias. Darkest brown regions reflect the fewest mental health providers and the most anti-gay sexuality bias.