Race as a Bundle of Sticks: Designs that Estimate Effects of Seemingly Immutable Characteristics

Maya Sen[1] and Omar Wasow[2]

[1]Harvard Kennedy School, Cambridge, Massachusetts 02138; email: maya_sen@hks.harvard.edu

[2]Department of Politics, Princeton University, Princeton, New Jersey 08544; email: owasow@princeton.edu

Annu. Rev. Polit. Sci. 2016. 19:499–522

The *Annual Review of Political Science* is online at polisci.annualreviews.org

This article's doi: 10.1146/annurev-polisci-032015-010015

Keywords

race, causality, research design, statistical methods

Abstract

Although understanding the role of race, ethnicity, and identity is central to political science, methodological debates persist about whether it is possible to estimate the effect of something immutable. At the heart of the debate is an older theoretical question: Is race best understood under an essentialist or constructivist framework? In contrast to the "immutable characteristics" or essentialist approach, we argue that race should be operationalized as a "bundle of sticks" that can be disaggregated into elements. With elements of race, causal claims may be possible using two designs: (*a*) studies that measure the effect of exposure to a racial cue and (*b*) studies that exploit within-group variation to measure the effect of some manipulable element. These designs can reconcile scholarship on race and causation and offer a clear framework for future research.

No causation without manipulation. (Holland 1986)

INTRODUCTION

Questions about group identity are fundamental to political science. Studies attempting to estimate effects of race and ethnicity, however, inevitably encounter methodological problems. Could a scientist conduct an experiment in which subjects were randomly assigned to be of different races? The simple answer—clearly not—has led many to warn against estimating the effects of "immutable characteristics" like race or ethnicity (Gelman & Hill 2007; Holland 1986, 2008; Winship & Morgan 1999).

More specifically, scholars have argued that race poses two challenges. First, any kind of treatment should be manipulable by a researcher—for example, by varying administration of a vaccine or enrollment in a job training program. Race, however, is commonly understood as an immutable characteristic. Second, race is "assigned" before most other variables; that is, people are typically categorized into one race or another from birth. Considering effects of race along with factors that follow birth, such as educational attainment or class, risks introducing post-treatment bias. Thus, making statements about the causal effect of race or race-based variables has been widely thought to be a misguided enterprise.[1]

Partly in response, some social scientists studying causal effects of race and ethnicity have adopted narrower experimental manipulations, such as varying the "racial soundingness" of a name on a resume, to approximate random assignment of seemingly immutable characteristics (Bertrand & Mullainathan 2004). Although these techniques help identify causal effects of something associated with race, they also introduce additional challenges of definition and measurement. Is race an immutable characteristic if elements of race can be manipulated? Are traits like "racial soundingness" the same as race? If not, how do those traits map to other aspects of race or to broader racial categories? At the heart of these methodological puzzles is an even older debate as to the nature of race. Is race immutable, as a primordialist or essentialist framework suggests? Or is a constructivist framework in which race is conceptualized as a complex, socially constructed identity with many mutable facets a more useful methodological starting point?

In this article, we address these questions and propose a new framework for studying the impact of race, ethnicity, and other seemingly immutable characteristics. Building on the work of both constructivist and quantitative scholars, we propose that, in experimental or empirical contexts, race should be understood as a composite variable or "bundle of sticks." Conceptualizing race and ethnicity in constructivist terms allows race to be disaggregated into constitutive elements, some of which can be manipulated experimentally or changed through other types of interventions. In many cases, this approach resolves the conflict between the potential outcomes framework of causal inference and seemingly immutable characteristics such as race, gender, and sexual

[1]Although race is often defined as a biological inheritance and ethnicity as a cultural inheritance, we use "race" and "race and ethnicity" interchangeably for four reasons. First, many groups, such as US Hispanics, are categorized as a racial group in some contexts and as an ethnic group in others. Second, within social science, the term of choice often varies by region and subdiscipline. For example, the term ethnic minorities is used by many European social scientists to refer to groups that would be considered racial minorities within the United States. Similarly, many scholars of comparative politics use ethnicity as an umbrella term for categories that include race. Third, epigenetics suggests that biological, environmental, and cultural influences interact in ways that can make drawing clean lines between biology and culture challenging. Fourth, in many studies, culturally determined traits are used to estimate effects of race. See Chandra (2006) for an overview of the challenges associated with defining and classifying ethnic identity.

orientation.[2] This approach is also useful for research focused on descriptive, observational, or correlational analyses. Thinking about race as having constituent parts can clarify what precisely is being estimated when scholars attempt to understand how race and ethnicity operate in the world. Our approach sheds light on the mechanisms at play and illuminates paths for potential policy interventions.

We illustrate this way of thinking about race by delineating two kinds of research designs: (*a*) studies that measure the effect of exposing an individual or institution to some racial or ethnic signal and (*b*) studies that attempt to measure the effect of some manipulable element of race that varies within a single group.[3] In short, our approach reconciles race and causation for many types of research and unifies a diverse body of past research into two coherent methods that can be applied to future scholarship.

This article proceeds as follows. First, we review theories of race developed by existing scholarship. We then briefly explain the potential outcomes framework, lay out the key problems of making causal inferences within the "immutable characteristics" framework, and show how theorizing and operationalizing race differently can resolve many of these problems. Finally, we tie these threads together into a cohesive framework that highlights two research designs: exposure studies and within-group studies. Throughout, we point to successful social science research to clarify how race-based variables can—and cannot—be used by applied researchers working to extract causal inferences from experimental and observational studies.

THEORIES OF RACE

How race is defined determines how it can be operationalized in empirical or quantitative research. Two theories of race have dominated prior scholarship: essentialism and constructivism. Essentialism tends to view race in largely biological terms and to categorize populations by regions of ancestry and phenotype. The concept may have arisen from 15th-century Europeans' efforts to rationalize slavery and colonialism (Zuberi 2001) and developed as 18th-century naturalists sought to classify populations from around the world (James 2011). From that work emerged the idea that members of groups shared "'essence(s)' that are inherent, innate, or otherwise fixed" (Morning 2011, p. 12), also described as "beliefs that a given social category is discrete, uniform, informative, . . . natural, immutable, stable, inherent, exclusive, and necessary" (Haslam et al. 2000; Morning 2011, p. 12). In the late 18th century, social Darwinists and eugenicists adopted ideas of race and advocated concepts of racial hierarchy that profoundly influenced how race was understood to work across science, politics, and society at large. In the 19th and 20th centuries, movements for and against white supremacy, as well as other forms of race-based nationalism, generated many of the inter- and intranational conflicts that defined those centuries (Du Bois [1903] 2007).

Although explicit arguments for racial hierarchy have moved from the mainstream of society to the margins, racial essentialism continues to inform how both lay people and scientists understand group differences (Mendelberg 2001, Morning 2011). Further, scholarly debates continue over how race and genetics determine intelligence, health, and other major life outcomes (Devlin

[2]This approach complements but is distinct from the concept of "intersectionality" (Crenshaw 1991). Whereas intersectionality examines the joint effect of multiple identities, e.g., the intersection of race and gender, the "bundle of sticks" approach seeks to disaggregate broad categories such as race into their narrower constitutive elements.

[3]Although we focus on race and ethnicity, much of this analysis and both research designs could also be used to estimate effects of other seemingly immutable characteristics (see, e.g., Boker et al. 2011).

1997, Duster 2005, Hernstein & Murray 1994). Some contemporary genetic research supports the idea that people with similar geographic ancestry also share clusters of common genes that correspond roughly to modern racial categories (Blank et al. 2004, Kitcher 2007; for a more thorough treatment, see James 2011).

The second theory of race emphasizes the weak scientific basis for racial categories and argues that race is best understood as a social construction (Appiah 1985, Omi & Winant 1994, Zuckerman 1990). In contrast to essentialism, the constructivist approach holds that distinctions between so-called races and the importance ascribed to various genetic or phenotypic traits are the products of social forces including cultural, historical, ideological, geographical, and legal influences (Holland 2008, Junn & Masuoka 2008, López 1994, Loury 2002, Rutter & Tienda 2005). How societies categorize difference typically reflects social structures that reinforce group-based hierarchy (Omi & Winant 1994, Sidanius & Pratto 2001).

Although most popular conceptions of race tend toward the essentialist, a considerable body of work suggests that a constructivist theory better fits how race actually operates in the world. For example, a 1974 US federal ad hoc committee on racial and ethnic definitions struggled with how to categorize people of South Asian ancestry who, earlier in the century, were categorized as Hindus or Hindoos (Hochschild & Powell 2008). The ad hoc committee initially recommended a designation of White/Caucasian but then selected the classification of Asian or Pacific Islanders (Nobles 2000). Penner & Saperstein (2008) find that in a 19-year survey of 12,686 Americans, 20% of the sample changed race in terms of either self-identification or classification by interviewers. Numerous other examples arise in the changing conceptions of what constitutes an interracial marriage or how children of mixed-race unions should be categorized (Kennedy 2012).

Many social scientists assume constructivism has become the standard academic approach, but research suggests otherwise. A 2011 survey of faculty in anthropology and biology departments across public and private universities found that only among more elite anthropology departments did a majority of the faculty define race as socially constructed (Morning 2011, figure 8, p. 182). Among biology faculty, race was defined as socially constructed by <15% of the sample from state universities and <40% of the Ivy League faculty. Similarly, 65% of college students defined race solely as biological. Among biology majors, 83% defined race as biological and 0% as a social construct (Morning 2011, table 4, p. 175). A more recent cross-discipline examination of scholarly articles finds that those in the hard sciences are more likely to express enthusiasm or optimism for genetics and genomics technology than are those in the social sciences and humanities (Hochschild & Sen 2015); one reason, the authors posit, might be that anthropologists and humanists adhere to a broader constructivist world view, which cautions against exclusive predictive emphasis on genetic information.

Turning to political science, most scholarship on race and causation has implicitly relied on essentialist ideas. Within comparative politics, many studies include dummy variables representing different "racial" or "ethnic" groups; in American politics or public opinion research, many studies include race as a set of dummy variables for analyzing differences among individual respondents. Thus, most research has assumed race to be an immutable characteristic inconsistent with the demands of causal inference.[4] Some causal inference scholarship has taken a more constructivist approach, but the methodological significance has, to date, remained undeveloped. Holland (2008), for example, defines race as a "socially determined construction with complex biological associations" (p. 95) but does not pursue the methodological implications. In the sections that

[4]Why essentialist ideas have predominated is unclear. Zuberi & Bonilla-Silva (2008) argue this is, in part, the product of the particular racial and ethnic experiences of those conceptualizing race as an immutable characteristic.

follow, we build on the concerns about immutable characteristics but operationalize race within the constructivist framework and show that estimating effects of race and ethnicity need not be ambiguous nor incompatible with causal inference.

CAUSAL INFERENCE AND POTENTIAL OUTCOMES

Does a vaccine cause people to live longer? Is a worker training program effective in helping people find employment? At its core, a causal inquiry involves unpacking the effect of some treatment on some outcome in which there is (*a*) a unit of analysis, (*b*) a manipulable treatment, and (*c*) a specific outcome. [The literature on the potential outcomes framework is voluminous—e.g., Angrist et al. (1996), Holland (1986), Splawa-Neyman et al. (1990), Rubin (1974, 2005)—and we attempt only a bare-bones introduction.] The fundamental problem of causal inference is, however, that we can never observe the difference between these two potential outcomes for any individual unit (Holland 1986, Rubin 1978). No single unit can receive both the treatment and the control at the same time. This problem extends to all kinds of inquiries, but it becomes particularly vexing with seemingly immutable characteristics.

In lieu of trying to estimate an unobservable true treatment effect, those interested in making causal inferences usually estimate some version of the average treatment effect, which is the difference between the mean outcome in treated and control populations. An obvious problem is, however, that differences in the outcome variable could be due to inherent differences between the treated and control populations, a problem some call selection bias (Angrist & Pischke 2009). For example, we should not be surprised to see that workers who have signed up for a worker training program are more successful in getting jobs—but we also should not be surprised that they are more ambitious and better educated than nontrained workers.

To get at a satisfactory estimate of the average treatment effect, we would like our treatment and control groups to be similar across all background variables that could affect both the probability of receiving treatment and the eventual outcome, so that the only difference between the two groups is that one received the treatment and the other did not. Many empirical efforts are geared toward trying to satisfying this ignorability requirement—that is, to make the treated and control populations similar enough that the treatment regime can be assumed to be random. By far the easiest course is simply to assign the treatment randomly, such as in a randomized experiment (for a more general discussion, see Holland 1986, Imai et al. 2008). However, because randomization is rarely an option for political scientists, and especially elusive for those studying race or ethnicity, researchers have turned to a variety of methods, like instrumental variables or controlling for observed variables, to satisfy the ignorability assumption and infer causal effects with observational data (Dehejia & Wahba 2002, Sekhon 2009).

CHALLENGES OF CAUSAL INFERENCE WITH RACE

The literature has identified two key problems within the context of race and potential outcomes. First, race is resistant to manipulation; second, because race is generally understood to be "assigned" at conception, the characteristics for which most social scientists control (education, income, etc.) occur after the treatment is assigned and therefore have the potential to introduce post-treatment bias (Greiner & Rubin 2010). In addition, we introduce a third problem: Race is unstable. By this we mean both that (*a*) across groups and time, the boundaries defining racial and ethnic categories are in flux and (*b*) within groups, there is substantial variation. This complexity may violate the requirement that a treatment should be comparable across observations.

Problem 1: Race Cannot Be Manipulated

Making causal inferences usually demands a neatly defined, manipulable treatment variable. Holland (1986), for example, famously admonishes "No causation without manipulation," meaning that all pertinent potential outcomes must be defined in principle in order to make causal estimates possible in practice. Further, to define all potential outcomes, one must be able to conceptualize an experimental analogy that would lead to the possible outcomes. In other words, as Holland (1986, p. 954) puts it, "causes are only those things that could, in principle, be treatments in experiments." The importance of a manipulable treatment is affirmed by many scholars (e.g., Cook & Campbell 1979, p. 36; Gelman & Hill 2007, p. 186; Pearl 2000).[5]

In an essentialist framework, however, race is resistant to manipulation or intervention, making it difficult to imagine appropriate counterfactuals. Imbens & Rubin (2010) refer to race and gender as "currently" immutable characteristics, as future scientific innovations may dramatically ease the effort required to change to seemingly fixed aspects of these characteristics. We can imagine how someone lives as an African-American; much more difficult is imagining an experiment or intervention that could manipulate the person's race (and only the person's race) so we could check its effect on some outcome. Not only is randomization beyond our reach, but even conceptualizing an ideal experiment or policy intervention is extremely difficult. As noted by Holland (1986, p. 946): "For causal inference, it is critical that each unit be potentially exposable to any of the causes. As an example, the schooling a student receives can be a cause, in our sense, of the student's performance on a test, whereas the student's race or gender cannot." Ultimately, as Angrist & Pischke (2009) point out, research questions for which there are no experimental analogies (even hypothetical ones, in a world with unlimited time and research budgets and omniscient powers) are fundamentally unidentified questions.

Problem 2: If Race Is the Treatment, Everything Is Post-Treatment

A second problem with conceptualizing potential outcomes is that a person's race, according to the "immutable characteristics" approach, is "assigned" at conception. Thus, the background covariates that social scientists usually control for or match on, such as education, income, and age, are determined after a person's race is assigned. Taking into account things that happen after the treatment happens has the potential of introducing post-treatment bias, a pervasive problem within observational social science research (King et al. 1994, Rosenbaum 2002).

To use a common example, suppose we are interested in the causal effect of smoking on death and have a population of randomly assigned smokers and randomly assigned nonsmokers. Should we control for lung cancer in the final analysis? No, because lung cancer is not only highly predictive of death but is also a direct consequence of smoking. If we controlled for lung cancer, the effect of smoking on death would be biased downward by the fact that we have controlled for its primary consequence. Race is obviously different from smoking, but the post-treatment issue applies with equal or greater force: Race deeply affects how a person is raised and educated, what employment opportunities he or she will have, and what cultural and social attitudes he or she will hold. Race, in other words, affects nearly every socioeconomic variable typically included in standard regression analyses, including ones meant to detect mediating patterns. Including any of

[5] Although a rich and varied literature (scholarly as well as popular) has developed around how multiracial people self-identify, these experiences represent a third kind of "treatment"—a mixed-race or racially ambiguous treatment (Faulkner [1932] 1990, Gates 1997, Griffin [1962] 1996, Halsell 1969, Hochschild & Weaver 2010, Kim & Lee 2001, Schuyler [1931] 1971).

these attributes could affect estimates of the causal effect of race, and not necessarily in a purely conservative direction. Thus, the existing practice of interpreting the residual impact of race is at best poorly conceptualized and at worst introduces serious bias.

Although perhaps unsatisfactory to many applied researchers, the most appropriate initial approach is to drop any post-treatment variables from an analysis (Gelman & Hill 2007, King 1991, King et al. 1994, King & Zeng 2006). In this context, any factor, attribute, personality trait, or personal or professional experience that could potentially be a consequence of race should be dropped. For example, if we were studying the effect of race on employment, we would not control for age, education level, income, criminal record, zip code, or health status, all of which could be impacted by the subject's race. The right-hand side of a regression would simply include race and possibly sex.[6] We note that this strategy implies that the researcher is interested in the *total* effect of race—which might not be satisfying to researchers or those unfamiliar with the causal literature (VanderWeele & Hernán 2012). However, there may be instances where the researcher is interested in the effects of constitutive components of race; we discuss this case below. This kind of research design still also fails to address the critique above that experimental analogies are undefined.

Even aside from the post-treatment issue, we note two further problems with controlling for race-related covariates: the common support problem and multicollinearity. The common support problem arises when researchers include attributes that vary according to race (e.g., welfare status, participation in programs such as Head Start, diseases such as Tay-Sachs and sickle cell anemia). Because these traits are highly clustered within certain groups, it becomes difficult to find cross-race comparisons. For example, finding a sizable group of whites who have sickle cell anemia would be challenging (Thomas & Zarda 2010). Collinearity becomes a problem when variables or effects vary so closely with race as to result in (the most extreme case) unconverged calculations of point estimates. The lack of variance in the background variables may also result in small changes having a large impact on the coefficient estimates—thus, standard errors may be large and lead researchers to assume no treatment effects when treatment effects do in fact exist.

Problem 3: Race Is Unstable

Building on the work of constructivists, we propose a third issue that is largely unaddressed by methodologists: Race is unstable and can vary significantly across treatments, observations, and time (Lee 2008, Abdelal et al. 2009). The category "Latino," for example, includes first-generation Mexican-Americans from Los Angeles and fourth-generation Puerto Ricans from the Bronx. In one analysis of census data, between 2000 and 2010 nearly 10 million respondents changed their self-identified race and/or Hispanic origin (Liebler et al. 2014). In quantitative terms, "no two measures of race will capture the same information" (Saperstein 2006, p. 57). This is true both across different studies and within the same study. For example, Bertrand & Mullainathan (2004) report that the treatment of receiving the name "Ebony" on a resume produced significantly different outcomes from receiving the name "Aisha," even though both are ostensibly the same treatment—a distinctively black name.

The dynamic and variable nature of race and ethnicity extends well beyond names. Bertrand & Mullainathan (2004) mention that they considered "other potential manipulations of race, such

[6]Sex, which is also assigned at conception, is one of the few standard control variables that are not post-treatment. We note, however, that some evidence suggests sex ratios can vary by latitude, religion, ethnicity, and other factors collinear with race (Guttentag & Secord 1983, Navara 2009). Other possibly pretreatment factors (e.g., genotype) are discussed by VanderWeele & Hernán (2012).

as affiliation with a minority group," but opted against out of a concern that "such affiliations might convey more than race" (p. 995, footnote 17). In other studies, subtle changes in cues such as wording in surveys or clothing in images resulted in significant differences in how race or ethnicity operated as treatments (Freeman et al. 2011, Sniderman & Piazza 1993). Research that fails to recognize this variability may violate the stable unit treatment value assumption (SUTVA), which requires that the treatment status of any unit does not interfere with the outcomes of other units and that the treatment "dosage" is comparable across all units. Forcing something as complicated as race into simple binary or categorical variables potentially complicates what we mean by a treatment. This is a problem not only for research designs focused on causal inference but also for those pursuing noncausal inquiries.

RESOLVING PROBLEMS WITH RACE AS A BUNDLE OF STICKS

Although the problems of causal inference with race can never be fully solved, in some instances they can be circumvented by theorizing race differently and using an appropriate research design. With regard to theory, we encourage empirical scholars to move away from defining race through an essentialist frame. For many questions, a constructivist frame is not only a better fit for the data but can also resolve problems of instability, manipulability, and post-treatment bias.

The problem of race as a potentially unstable treatment can be addressed, in part, by exploiting the constructivist observation that race is rarely if ever a single, uniform entity. As scholars in race and ethnic politics, sociology, anthropology, and critical race theory have emphasized repeatedly, racial categories are the product of a complex fusion of factors including societal values, skin color, cultural traits, physical attributes, diet, region of ancestry, institutional power relationships, and education. In other words, race is an aggregate of many components; metaphorically, it is a bundle of sticks (**Figure 1**). In contrast to the "immutable characteristics" approach, we argue that race is most accurately understood as a composite measure that can, in some cases, be disaggregated into constitutive elements. Elements of race that are strongly identified with or highly collinear

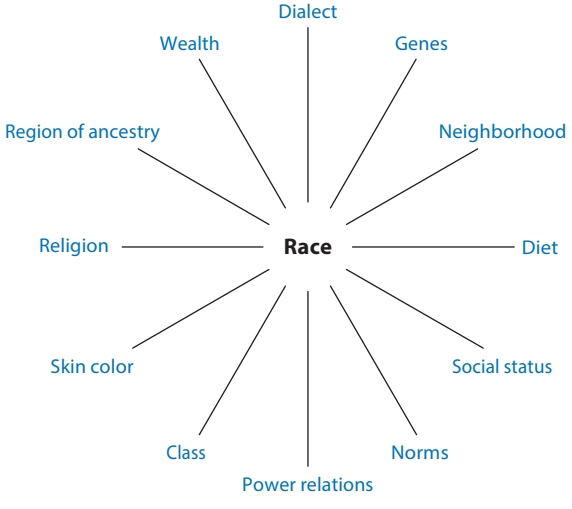

Figure 1

Some characteristics associated with race and ethnicity.

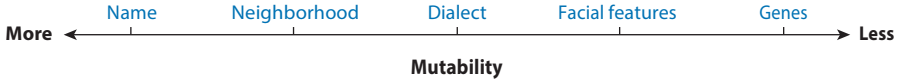

Figure 2

Hypothetical mutability of characteristics associated with race and ethnicity.

with the particular racial or ethnic category can be thought of as constitutive or what make the composite of race and ethnicity meaningful in the world.

This is not only a much more tractable enterprise but also has the advantage of solving one of the most persistent problems associated with studying race or ethnicity: the difficulty of knowing what exactly is being estimated. A randomized medical trial, for example, that incorporated multiple changes in a diet (e.g., the Mediterranean diet) would be unable to distinguish which elements of the dietary intervention were therapeutic. Only by isolating a single change, e.g., supplementing omega-3 fatty acids, could a specific effect be identified. Most causal (or even most descriptive) estimands fail to capture the entire bundle of attributes that constitute race and instead capture some component.

To clarify this approach, we analogize to another commonly used composite variable, socioeconomic status (SES). SES is composed of family income, educational attainment, occupation, and other measures. Given its composite nature, experimentally manipulating all the elements of SES simultaneously would be difficult. Likewise, it would be problematic to make causal claims with any design that compared people with sharply different SES. We could, however, assess the causal effects of manipulating one element of SES, such as education, within a population of similarly situated subjects. By definition, measures of educational attainment and SES are distinct; but, also by definition, any change to the former will have a downstream effect on the latter. Hence, understanding an effect of education, all else held constant, will help explain an important part of the effect of SES. Similarly, once race is operationalized as a composite variable, estimating the effect of a substantive and constitutive element of race helps explain how race works.[7]

Once race is operationalized as a disaggregable composite variable rather than a monolithic, homogenous entity, the problem of manipulability can be resolved by identifying an element of race that is relevant to the research question at hand and can be manipulated in at least one of two ways. First, many seemingly "immutable characteristics," once disaggregated, are manipulable in the context of experiments. In audit studies, for example, researchers can send confederates into the field to apply for employment and randomly assign the job applicants to be from different racial categories. Similarly, in lab and field experiments, researchers can manipulate media with auditory or visual cues about otherwise hard-to-modify elements of race (**Figure 2**).

Second, many elements of race are, in fact, mutable. **Figure 2** presents a hypothetical continuum of features that are associated with race but exhibit varying degrees of mutability. Facial features—such as the shape of one's eyes or the contours of one's nose—are fairly immutable, possibly changed through plastic surgery but certainly not something researchers could easily manipulate in a study or policy intervention.[8] In many experimental contexts, such traits are less useful as they present the same conundrums identified by the "immutable characteristics" framework.

[7]One important difference between SES and race is that the former tends to be coded as a continuous variable and the latter as a discrete variable. As such, manipulations of elements of race may produce "lumpier" effects in things like racial categorization. Even within a discrete coding of race, however, it may be possible to use continuous measures of factors such as degrees of identification with a group (see, e.g., Knowles & Peng 2005).

[8]The boom in ethnic-oriented plastic surgery might present some interesting, if far-flung, experimental possibilities (Dolnick 2011, O'Connor 2014).

However, traits that are highly collinear with race and mutable are often well suited to causal inference. They are also more likely to be the product of social and environmental forces. For example, a large literature in gender studies distinguishes between "sex" and "gender": "Sex" is defined as biological and anatomical whereas "gender" is defined as the product of psychological, social, institutional, and cultural forces (see, e.g., Deaux 1985, Htun 2005, West & Zimmerman 1987). Similarly, where appropriate, we suggest scholars of race and ethnicity consider distinguishing between less mutable, typically biologically ascribed correlates of race and more mutable, typically socially or environmentally assigned aspects of race (with the understanding that such categories can never be cleanly delineated). Environmental interactions are also important to consider, as many seemingly immutable biologically inherited characteristics, such as skin color or alcohol flush reaction, are responsive to triggers such as sun exposure or drinking wine.

Finally, the problem of post-treatment bias can be resolved in cases where constitutive elements of race are assigned after conception or remain manipulable after conception. Newborn infants, for example, exhibit no preference for faces from their own racial or ethnic group, but three-month-old infants do (Kelly et al. 2005). Bar-Haim et al. (2006) find that this early encoding of own-group visual preferences can be attenuated by exposure to individuals from another race. Similarly, birth weight can vary significantly by race, but evidence from twin studies and other natural experiments suggests that a variety of manipulable factors, such as maternal access to food stamps, can positively influence intrauterine nutrition, birth weight, neonatal mortality, adult schooling attainment, height, and, for lower-birth-weight babies, labor market payoffs (Almond et al. 2011, Behrman & Rosenzweig 2004, Conley & Strully 2012). Research in life course epidemiology and epigenetics further suggests that many constitutive elements of race are assigned by social and environmental forces after conception or birth. Maternal stress, early-life undernutrition, and other early-life forces become "embodied" and durable points of differentiation across adult populations defined by racial and ethnic categories (Ben-Shlomo & Kuh 2002, Kuzawa & Sweet 2009).

A variety of adult life experiences can also shape racial identification and categorization. Living in the suburbs, receiving welfare, or being incarcerated can influence how people self-categorize by race and are perceived racially (Penner & Saperstein 2008, Saperstein & Penner 2010). How people die also influences racial classification: Noymer et al. (2011) find that on death certificates, victims of homicide are more likely to be classified as black and people who die of cirrhosis of the liver are more likely to be classified as American Indian, even when controlling for a separate racial classification offered by the decedents' next of kin. Traits such as language and dialect are also highly collinear with racial and ethnic background but are mutable and assigned postconception. Purnell et al. (1999) make telephone calls to landlords and find significant "linguistic profiling" and racial discrimination against potential tenants on the basis of dialect.

In short, when operationalized as a composite variable, race is disaggregable, some "sticks" are manipulable, and the whole bundle is not automatically assigned at conception. In addition, the more mutable characteristics represent attributes that could serve as plausible interventions, including potential policy interventions; that is, we cannot conceptualize how policy actors would intervene in terms of assigning people to one race or another under an essentialist framework, but we can certainly think about meaningful, plausible policy prescriptions whereby subjects from different racial or ethnic backgrounds are assigned different names, neighborhoods, income transfers, or diets. Not only does our approach enable these important inquiries, but it does so without running afoul of the potential outcomes framework. **Table 1** summarizes how race is operationalized within both the "immutable characteristics" and the "bundle of sticks" frameworks.

Table 1 Summary of the "immutable characteristics" versus "bundle of sticks" approaches to operationalizing race

Operationalization of race	"Immutable characteristics"	"Bundle of sticks"
Underlying theory	Essentialist	Constructivist
Race manipulable?	No, race is an immutable characteristic	Yes, race contains mutable and manipulable elements
Always post-treatment bias?	Yes, race is assigned at conception	No, some constitutive elements of race are assigned after conception
Race unstable?	No, race is homogenous and measurable	Yes, race demands disaggregation
Measurement?	Race is typically coded as a binary or categorical variable	Race is a composite variable in which an element of race is the key variable and determines coding

In addition to rethinking how race is operationalized, we encourage scholars to consider whether the question being investigated can be addressed by one of the two research designs we discuss in the remainder of this article. In the first design, an element of race operates as a cue or signal that generates some reaction. In the second design, an element of race exhibits within-group variation and partly explains how the larger composite of race shapes life outcomes. We call the first type an exposure design and the second a within-group design. Exposure studies are ideal for studying discrimination or implicit bias, as an element of race typically acts as a proxy when researchers attempt to estimate an effect of the larger bundle of race. For example, names often act as a proxy for traits associated with racial or ethnic groups. In within-group studies, an element of race is identified to estimate the effect of one part or "stick" in the larger whole. As an example, we might study the role of birth weight as a contributor to racially disparate academic achievement. Both approaches also suggest more meaningful and tractable policy interventions than, say, attempting to understand the effect of race as a whole.

RESEARCH DESIGN 1: EXPOSURE STUDIES

Exposure to a racial cue or signal conveys information about race to a subject. Exposure studies have been described as those that look at the effects of "perceived race" (Greiner & Rubin 2010) and discrimination (VanderWeele & Hernán 2012). We use different terminology and draw different analogies, but the research designs we suggest here are comparable.

We move away from the "perceived race" and discrimination language for three reasons. First, we think the best way to think about the treatment in exposure studies is not as perception but instead as a signal about race. After all, in an experimental context, the researcher can manipulate the signal to which the subject is exposed but not what the subject actually perceives. Second, perceived race is rarely observed; perception occurs within the subject's mind and is generally opaque to researchers.[9] As such, focusing on exposure to a racial signal rather than perception of race is preferable. Finally, not all studies involving exposure to a racial cue involve discrimination as conventionally understood. Studies of "stereotype threat," for example, have exposed female and minority students to racial and gender cues prior to taking an exam (Steele 1997). Rather than triggering discrimination by some external source, the cues trigger internal anxiety about

[9]Many experiments pretest treatments and/or run post-treatment manipulation checks, but, even then, much of what subjects perceive remains unobserved.

confirming negative stereotypes.[10] We prefer to categorize this design by the method of treatment and to be agnostic about the particular context or outcomes of the intervention.

In this research design, (*a*) one or more elements of race is identified as a relevant cue; (*b*) subjects are treated by exposure to the racial cue; (*c*) the unit of analysis is the individual or institution being exposed. All three steps alleviate the problems of race and causality. The research design begins with well-defined potential outcomes, is operationalized via a clean experiment (or a clean experimental analogy), and has a precise moment of treatment. Through a proxy for race as a whole, a causal impact of race and ethnicity is identified, alleviating the problems of manipulability, instability, and post-treatment bias.

Experimental Exposure Studies

Studies across the social sciences have used exposure to a racial or ethnic signal as a key feature of the experimental design. In sociology and economics, audit and correspondence studies have been used to measure racial and other forms of discrimination, typically in field experiments. Audit studies usually involve confederates or actors hired by researchers who are then randomly sent out to the field. Pager (2003), for example, sent men to apply for working-class jobs and randomly assigned the applicants by race and other attributes. Partly in response to critiques about potential bias introduced by the confederates, correspondence studies, in which matched human applicants were replaced with matched pairs of "paper" applicants, have become more common (Heckman 1998, Heckman & Siegelman 1993; for a good overview of the literature, critiques, and methods, see Pager 2007). In political science, Butler & Broockman (2011) and Broockman (2013) used distinctively black and white names in putative "constituent" emails to legislators.

In sociology and political science, survey experiments with racial signals are now regularly used to estimate effects of race. These experiments typically manipulate survey questions or media, such as newspaper articles or political campaign ads, to estimate how randomly assigned racial cues influence attitudes and behavior. Sniderman & Piazza (1993), for example, leverage question order to find that the "mere mention" of race-based affirmative action to white survey respondents provokes more negative feelings toward blacks. A robust public opinion literature exploits some variant of the "exposure to a racial signal" design to estimate causal effects of race (Gilens 1996, Huber & Lapinski 2006, Miller & Krosnick 2000, Tesler 2012, White 2007). Mendelberg (2001) and Gilliam & Iyengar (2000), for example, create simulated television news experiments to assess how racial cues might prime racial attitudes among white voters. Similarly, Valentino et al. (2002) test whether subtle racial cues in campaign advertisements prime racial attitudes and candidate preference. Framing experiments by Bobo & Johnson (2004) use survey questions about criminal

[10] Some scholars suggest that what we describe as an effect of race is more accurately called an effect of racism (R. Kramer, Twitter discussion with authors, **https://twitter.com/rory_kramer/status/503564340226973696** and **https://twitter.com/rory_kramer/status/503564598726111232**, Aug. 24, 2014). We reject this suggestion for three reasons. First, within this framework the researcher typically manipulates a cue but not the context in which the cue is received (e.g., a society with high rates of bigotry). We agree that outcomes of interest are often the joint effect of the cue and the social context but find it conceptually more useful and clear to focus on the specific variation identified by the scholar. Second, although we agree many outcomes of interest are directly influenced by bigotry, we find it reductive to assume all effects of race are products of racism. Many aspects of a racial or ethnic experience (e.g., traditions, cuisines, norms, etc.) precede the rise of modern bigotry. In addition, some important population-level differences, such as health disparities, are likely influenced by regions of ancestry rather than discrimination. For example, the fact that white Americans get skin cancer at much higher rates than African-Americans is unlikely to a be a product of racial bias. Third, some forms of bias (e.g., a baby's preference for faces from her own racial group) would not typically be understood as racism. Other forms of in-group bias, such as assortative mating or a broader human tendency toward homophily, may also operate differently than racism. Should scholars prefer to describe these phenomena as effects of racism, the basic framework we outline remains the same.

justice to estimate how different racial cues shape the "taste for punishment." Gay & Hochschild (2010) conduct a survey experiment to assess the breadth of feelings of "linked fate" by varying racial, gender, and other identity cues in question content and ordering (Dawson 1994).

A growing body of research in political science evaluates the effects of racial cues on voting behavior. Green (2004), working with the NAACP National Voter Fund, evaluates whether phone calls from other African-Americans and direct mail crafted to appeal to the concerns of African-Americans increased voter turnout. Enos (2011) tests a subtle form of racial threat by mailing voters information about proximate outgroup voting rates. Valenzuela & Michelson (2011) conduct a get-out-the-vote experiment in which Latino-surnamed voters receive calls that cue either ethnic or national group identities. Language also matters for political mobilization (Bedolla & Michelson 2012). Abrajano & Panagopoulos (2011) find significant effects of English- versus Spanish-language appeals in a get-out-the-vote campaign targeting Latinos.

Studies in psychology, and related fields such as political psychology and behavioral economics, suggest additional types of exposure to a racial signal. Steele (1997) identifies how internalized stereotypes affect women and racial minorities. The Implicit Association Test (IAT) developed by Greenwald et al. (1998) measures response latencies when subjects are given the assignment to quickly categorize stimuli, often words and images with racial cues, into pairs of categories. Kurzban et al. (2001) expose subjects to images of a hypothetical cross-race conversation and use errors in recall to assess if and how race is encoded in memory.

Although these studies are able to cleanly identify effects, we note several possible sources of confusion as to what exactly is being identified. Racial and ethnic cues can generate meaningful effects only when they trigger thoughts that subjects associate with a particular group in a particular context. Consequently, racial signals should always be understood to operate as a joint effect of the cue and the social, political, and historical context in which the experiment occurs. Failure to distinguish between the cue, the context, and the joint effect can lead to at least three issues.

First, studies may overstate claims about identifying the causal impact of race when, in fact, only an element of race has been experimentally manipulated. Scholars should be clear about which constitutive component of race or ethnicity is serving as the treatment. In addition, to make claims about a broader effect of race, scholars should state their assumptions about the link between the element of race or ethnicity being studied and the identity category as a whole (e.g., dialect serves as a proxy for race as a whole). Where possible, researchers should also pretest the link between the cue and how subjects interpret the signal in terms of identity.

Second, some studies are careful to report the effect only in terms of an element of race (e.g., "racial soundingness of a name") and fail to convey that the narrow cue likely exhibits powerful effects by triggering associations with race as a whole. Here, precision in describing the treatment can lead scholars to understate or even overlook the fact that the race cue only works as a joint effect with other associations such as racist beliefs.

Finally, even when a seemingly narrow element of race has been employed to identify broader effects of race, the cue may still encode other information or "sticks" that confound interpretation. This problem can arise when conceiving of racial categories as coherent, homogenous entities. As noted above, Bertrand & Mullainathan's (2004) pathbreaking study shows that resumes with the first name "Ebony" elicit calls from potential employers 9.6% of the time whereas otherwise identical resumes with "Aisha" have a callback rate of 2.2%. The authors acknowledge "significant variation in callback rates by name" (pp. 1,008–9) for African-American females but the possible heterogeneity in the "black" treatment remains unexplained within a binary or categorical model of race.

Although Bertrand & Mullainathan did pretest the names as racial cues, their results suggest the pretest did not capture the full range of information conveyed by seemingly similar "black"

names. As we emphasize below, these issues can often be resolved through greater attention to what specifically constitutes the treatment and which component of race is being captured.

Although scholars have long viewed audit and correspondence studies as related, we argue that *all* studies employing exposure to a racial or ethnic signal share a common experimental design. These studies exploit different techniques—from simulated avatars to scenarios in surveys—but the general approach is the same: randomly present a subject with information that differs only with respect to signals or cues about race or ethnicity. It is important to note that the treatment is never all traits associated with race (i.e., the whole bundle of sticks) but only an element of race that serves as a proxy for the bundle. Moreover, the meaning ascribed by subjects to the bundle depends heavily on the combined effect of the cue and the context in which the cue is observed.

Observational Exposure Studies

It is possible to import this research design to a wide variety of observational contexts involving how third parties react once they are exposed to racial signals and cues. Greiner & Rubin (2010), for example, investigate how juries react to Hispanic versus non-Hispanic death penalty defendants, and Wasow (2012) explores how white voters respond to exposure to protests by blacks that escalate to violence. In these instances, the interest lies in understanding how exposure to a racial signal changes or informs opinions, behaviors, or attitudes. Researchers working with observational data can structure their analyses to approximate an experimental exposure design. This design is often ideal for testing implicit bias or racial discrimination (Greiner & Rubin 2010, VanderWeele & Hernán 2012).

Researchers inferring causal effects from observational data must be aware of two attendant issues. First, using observational data means that researchers lack the ability to manipulate the racial cues and signals received by the subject. It is therefore necessary to use techniques such as matching or inclusion of control variables in a regression model such that the only observed difference between the treated and control groups is that they are exposed to distinct racial signals (including the possibility that one group receives no racial cue at all). This means that these research designs still must confront the possibility of unmeasured confounders—e.g., those factors that could correlate with race or ethnicity (and could affect the outcome) that are not captured by the set of covariates included in an analysis.

In theory, if all confounders are accounted for in a model, a reasonable assumption would be that the residual impact of race is the "causal effect" of race; that is, the effect of race not captured by the other covariates. In practice, this condition is never met, and we caution against interpreting the residual in this manner. Generally, it is impossible to know whether all unobserved variables have been included in a model. Moreover, once race is operationalized as a composite variable, what is commonly described as the residual effect of race or ethnicity should be understood as an estimate of the composite effect of all the unobserved elements of race (including possible interactions of any observed and/or unobserved terms).

For example, imagine a simple scenario in which a composite measure of race can be generated by using the variables in **Figure 2**. A regression model that included half of the variables as controls and a term for "race" would be estimating the joint effect of the other half of the variables. In many cases, if all relevant measures were truly accounted for in a model, the residual effect of race would approach zero and there would be little to no independent effect of race. In either case, there may be some or no evident residual effect of race, depending on how race is operationalized and on what other variables are included in the model in which the race term is used.[11]

[11] As there is no way to measure unobserved confounders, we note that sensitivity analyses are a useful way of at least estimating

Second, and perhaps more helpfully, the exposure design can lessen problems of post-treatment bias (Greiner & Rubin 2010) but requires researchers' vigilance. Suppose we are interested in whether a bank offers different interest rates to minority versus nonminority loan applicants. The ideal experiment would be to mimic an audit study and create identical loan applicants whose profiles differ only with regard to how they are categorized into racial groups. The "treatment" would be the loan officer's review of the application packet. Anything that happens before is solidly pretreatment and must be conditioned on; this would include anything that could potentially appear on an application for a loan. Anything that happens after the decision maker reaches a decision (e.g., extending additional credit, the size of the loan) would be post-treatment and should be dropped from the statistical model (Greiner & Rubin 2010). Again, drawing an analogy to the ideal exposure study is helpful in assessing which covariates could be construed as pretreatment and which could be construed as post-treatment.

This discussion can be boiled down to one key idea: When possible, conceptualizing an experiment or observational study as an "exposure to a racial signal" study greatly reduces both the theoretical and practical problems associated with making race-based causal inferences. Thus, applied researchers should think carefully about whether an exposure study could provide a well-suited analogy for their research questions and hypotheses.

RESEARCH DESIGN 2: WITHIN-GROUP STUDIES

Why is the lifetime risk of developing diabetes higher for Hispanics than for other groups? Why are certain ethnic groups overrepresented in rebel militias? Studies of such questions involve no clean treatment by exposure to a racial cue and no decision maker [in the terminology of Greiner & Rubin (2010)]. VanderWeele & Hernán (2012) refer to these studies as those focusing on discrepancies. This work is often attempting to understand how a part of race shapes the whole. For scholars working on these sorts of topics, the primary research interest—and the appropriate unit of analysis—lies in a particular racial or ethnic population itself. These studies are particularly problematic in terms of having ill-defined potential outcomes and post-treatment bias problems.

For such questions, we suggest a research design that exploits variation *within* a racial or ethnic group, not across groups. The within-group design disaggregates the bundle of sticks and singles out a specific constitutive element of race or ethnicity that can be manipulated in an experiment (or observed to vary) within a group. For within-group research designs, (*a*) one or more constitutive elements of race that exhibit within-group variation are identified as a treatment; (*b*) members of the group are assigned to the treatment and control conditions (or are observed to vary across the conditions); and (*c*) the units of analysis are the individual members of the group. As with the "exposure to racial cue" approach, these steps help mitigate the problems of race and causality. These steps also help isolate causal mechanisms and help scholars think more clearly about what could be more tractable and meaningful policy interventions.

For example, suppose we seek to understand disparate educational outcomes for black versus white youngsters. A naive analysis would be to regress educational outcomes on race, with the group of African-Americans as the treated group and whites as the control, possibly controlling for other relevant variables. For all the reasons cited above, however, a causal estimate based on this research design would be (*a*) fundamentally unidentified and (*b*) biased by any inclusion

their potential effects. These sensitivity tests place bounds on the size of the confounding that one would have to see among the treated group (e.g., the racial minority group) in order to render insignificant those effects that have been detected. Greiner & Rubin (2010) provide some useful examples, and Keele (2010) and Rosenbaum (2002) discuss the methodology.

of post-treatment variables.[12] Furthermore, such a naive regression would not isolate *why* black youngsters fare worse; after all, a statistically significant coefficient on the "black" variable would simply reveal that an education gap continues to exist. Last, such a design would probably not shed light on potential policy interventions to ameliorate such discrepancies.[13]

A better research design would start with the fact that race is composed of a variety of factors, and, rather than conceive of black youngsters as a treated group and white youngsters as the control, identify a trait that is (*a*) a possible explanation for the gap, (*b*) collinear with race, but not perfectly so, and (*c*) in theory, manipulable. One example might be neighborhood. With the long history of residential segregation in America, race and neighborhood are distinct but highly collinear. Neighborhood effects, through factors like variation in the quality of local schools or police, could plausibly explain part of the education gap, and neighborhood can be varied in ways that race cannot.

With this in mind, we can recast the study as a within-group analysis. We compare academic achievement by black youngsters from, say, high-poverty neighborhoods to similarly situated black youngsters in moderate-poverty neighborhoods. The Moving to Opportunity experiment, which incorporated random assignment of housing vouchers, offers one example of just such a design (Katz et al. 2001). Scarr et al. (1977) exploit variation in the degree of white ancestry within an African-American population and find that genes associated with Caucasian ancestry show no relationship to cognitive ability. By identifying meaningful within-group differences, scholars can narrow the causal mechanisms that explain disparate across-race outcomes.

This research design has several advantages over more naive cross-race regression approaches. First, limiting the unit of analysis to a single racial group and conceptualizing the treatment as being something that varies closely, but perhaps not exclusively, with race allows for experimental manipulation, in theory or practice. This not only permits us to avoid the critique that no well-defined potential outcomes exist, but also means that we can think of meaningful policy interventions to address race-related discrepancies. Second, because the alternative treatment may be "assigned" postbirth, this design also allows for the inclusion of all pretreatment variables (confounders), including such traits as mothers' education, health, nutrition, and early educational opportunities. In this regard, we could think of race or ethnicity as a confounding variable that can be controlled for or conditioned on.[14]

Third, with enough data, conditioning on race before moving to a causal analysis resolves the common support problem. It might be difficult to find a sufficient number of similarly situated individuals across racial groups, but focusing on within-race variation will often resolve this problem.

Experimental Within-Group Studies

A growing number of experimental studies, particularly in psychology, use the within-group approach. Walton & Cohen (2011), for example, randomly assign freshmen to receive a message

[12] A plausible way to rethink the research design in this example would be to take an SES variable as the treatment of interest and race or ethnicity as the pretreatment confounder. This would represent a different inquiry, albeit an interesting one.

[13] It might be tempting to try mediation analysis with these types of questions. For example, one could treat family income as a mediator. Identifying the effect of race on an outcome that passes through income would be difficult, however, without very strong assumptions. For example, to use traditional mediation analysis, race would have to be the only factor affecting income (Imai et al. 2011), an assumption that is clearly not met.

[14] For example, intervening on things like neighborhood, mothers' education, health, nutrition, and educational opportunities could have different effects across different groups—a kind of effects modification. Because the impact of the alternative treatment may vary according to subgroup, comparing the results between groups may also be useful. In our neighborhoods example, including comparisons with white children in the analysis might shed some light on these issues but would probably not help us make meaningful causal inferences.

that all college students struggle to fit in initially but can ultimately succeed. In this case, the constitutive element of race is an uncertain sense of belonging for stigmatized groups in school and work settings. Compared to the black control students, the black treated students exhibited substantial sustained academic improvements over their college careers and later reported being happier and healthier. Walton & Cohen (2011) also included a white comparison group and found that treated whites exhibited no significant differences from control-group whites. Put another way, uncertainty about social belonging in college appears to be sufficiently collinear with race as to be constitutive for African-Americans yet immaterial for whites. At the same time, feelings of social belonging are sufficiently malleable that a simple exercise lasting about 45 minutes could dramatically change outcomes for treated black students as compared to black controls.

In political science, Gay (2012) builds on the Moving To Opportunity experiment and investigates the role of high-poverty neighborhoods on voting. Gay finds that poor families offered vouchers to leave public housing vote at lower rates. Although Gay's analysis is not explicitly focused on explaining the effects of neighborhood as an element of race, the sample population in the study is nearly two-thirds black and nearly one-third Latino. As such, the analysis is implicitly a study of the role of neighborhood context and social dislocation as elements of race in minority turnout. Valenzuela & Michelson (2011) also explore the role of neighborhood context in a get-out-the-vote experiment by comparing the differential resonance of ethnic and national identity appeals across middle-class and working-class Latino communities.

Observational Within-Group Studies

Observational studies have also successfully leveraged components of race in order to extract surprising inferences. Sharkey (2010) exploits temporal variation in local homicides in Chicago to identify a significant neighborhood effect of proximity to violence on the cognitive performance of African-American children. Cutler et al. (2005) investigate why African-Americans suffer from higher rates of hypertension than do whites. By more closely examining black subpopulations, they demonstrate that blacks whose enslaved ancestors survived the Middle Passage across the Atlantic exhibit higher rates of salt sensitivity than do blacks whose ancestors were not enslaved (i.e., more recent African immigrants to the United States or the United Kingdom). A possible mechanism is that salt retention—a precursor to hypertension—enabled enslaved Africans to survive the deadly three-month sea voyage that constituted the Middle Passage. Thus, the appropriate "treatment" in this study was having ancestors who were subjected to the Middle Passage. Because no European-Americans were subjected to that voyage, the "treatment" is highly collinear with being African-American but not necessarily with being of African descent, a finding made clear only by within-group comparisons.

Nisbett & Cohen (1996) investigate high rates of violence among men in the American South. A typical cross-race approach, as is often used in fields such as health and education, might have compared rates of violence among white and black men. Owing to post-treatment bias, such comparisons are problematic if the researcher is attempting anything more than a descriptive analysis. Nisbett & Cohen, by contrast, exploit within-group variation among whites and avoid post-treatment bias pitfalls. Through both observational data and experiments, Nisbett & Cohen identify specific cultural traits that vary between Southern and Northern white men, which influence attitudes, physiology, and differential rates of violence.

As with other studies relying on observational data, researchers using within-group designs should consider experimental analogies. This point has been made by the causal inference and econometrics literatures but is particularly worthwhile for those specifically interested in race

Table 2 Overview of exposure and within-group research designs

	Exposure	Within-Group
Unit	Individuals or institutions, potentially from any group	Members of a particular group
Typical treatment	Racial cue or signal (e.g., include distinctively ethnic names on a resume)	Constitutive element of the composite of race (e.g., address anxiety about social belonging in college)
Role of element of race	One "stick" is a proxy for the bundle (e.g., in a phone call with a landlord, dialect signals many traits associated with race)	One "stick" explains part of the bundle (e.g., Middle Passage might partly explain high rates of hypertension among African-Americans)
Examples	Correspondence and audit studies Implicit Association Tests	Experimental manipulation of a constitutive psychological dimension of race Within-race matching

(Angrist & Pischke 2009). Keeping an eye on what the ideal experiment would look like (and what factors would or would not have to be controlled for) is essential for thinking clearly about potential identification strategies and problems. In addition, given the absence of randomization, researchers using within-group designs with observational data should use tools like matching and inclusion of pretreatment variables in regressions to address the ignorability assumption. **Table 2** summarizes key aspects of the exposure and within-group designs.

COMBINING EXPOSURE AND WITHIN-GROUP DESIGNS

It is possible in at least four cases to combine aspects of the exposure and within-group designs. First, some researchers may wish to use exposure designs solely with particular racial or ethnic subgroups. In this case, within-group variation is introduced by exposure to a racial cue, and the subject pool is narrowed to reduce heterogeneity among the observations. Lee & Pérez (2014), for example, evaluate language-of-interviewer effects on Latino public opinion and find substantial differences in respondents' attitudes and reporting of political facts.

Second, some researchers may be interested in how subjects respond to racial or ethnic cues in which at least some of the variation in signals occurs within rather than across groups. Adida et al. (2010), for example, apply to jobs with French employers in which resume names have been randomly assigned to signal a person of Senegalese and Christian background, Senegalese and Muslim background, or a "typical French republican" background with no religious affiliation. Hopkins (2015) exploits differences in immigrant skin tone, language, and accent to experimentally vary within-group racial cues in the context of a TV news segment. Both examples use an exposure design in which the cues involve race and traits that vary within race, like religion or accent. In this design, subjects—potentially of any background—are exposed to cues but the signals are not exclusively cross-racial or cross-ethnic.

Third, a combined design can be useful for assessing interaction effects between within-group traits and exposure to a cue. Valenzuela & Michelson's (2011) study, for example, compared receptivity to ethnic or national group identity cues across Latino subgroups. This design allows for an estimate of the joint effect of a within-group trait (in this case, the class characteristics of the neighborhood) with priming effects of exposure to a cue. Here, the unit of analysis is the same as that of a within-group design in which the subjects are members of a single group and in which variation of some constitutive element of the group is exploited for causal inference. In essence,

each subject receives two treatments (i.e., within-group neighborhood characteristics and a racial or ethnic cue), and this design allows for causal inference about the combined effect.

Fourth, scholars may wish to compare results of an exposure study both within and across groups. Such studies typically involve two racial or ethnic groups that each have a separate treatment and control subgroup. Walton & Cohen (2011), as mentioned above, create black treated, black control, white treated, and white control groups. The treatment is exposure to media and some simple exercises that are designed to address anxieties about social belonging. The results of the social belonging intervention—big benefits for treated black students and essentially no effect for whites—are discernible only by combining the exposure to a racial cue and two within-group designs.

TOWARD A UNIFIED FRAMEWORK FOR RACE AND CAUSALITY

In this article, we have proposed a new way of thinking about estimating causal effects of race and ethnicity. First, we argued that social scientists should reconsider how they theorize and operationalize race. As shown by Morning (2011), the debate between essentialists and constructivists is far from resolved. In contrast to essentialist or "immutable characteristics" approaches, we argue here that a "bundle of sticks" conception better represents how race and ethnicity operate in the world. Moreover, operationalizing race as composite and disaggregable is more amenable to causal inference. Immutable and manipulable need not be incompatible. For those social scientists already disaggregating race but lacking any theoretical framework, our approach clarifies the relationship between an element of race being studied and the whole bundle. Rather than simply assuming connections, scholars can state that a particular element of race is a part of the larger composite or they can explain that the element of race is serving as a proxy for the whole.

Second, we have generalized two research designs appropriate for investigating causal effects of seemingly immutable characteristics. The exposure design may be particularly appropriate for those studying public opinion, political behavior, implicit bias, stereotype threat, law, and public policy—fields in which questions of interest frequently involve how institutions or individuals view and interact with racial signals and cues. For research focusing on features of particular populations, we encourage consideration of within-group designs that exploit constitutive, varying, and manipulable elements of race. Even though some aspects of race may not lend themselves to manipulation, many highly collinear elements of race may be experimentally manipulated or observationally assessed. Many important questions and cases are beyond the scope of the approaches we present, and appropriate elements of race may not always be available. Nevertheless, some elements may vary closely with race, may not already be included in the analysis, and may explain a significant part of the bundle.

A final reason we recommend the "bundle of sticks" approach is that it forces researchers to consider exactly what is being captured by racial identification variables. The multifaceted nature of race and ethnicity suggests that when race is operationalized as a stable, homogenous entity (e.g., a simple dummy or categorical variable like "1" if white, "0" if nonwhite), any statistical association will typically offer little or no insight as to which elements are the key mechanisms of action—be it fear of an out-group, neighborhood effects, or some other factor. Also, just as it is difficult to imagine a way to assign race experimentally, it is difficult to translate research identifying simple racial or ethnic disparities into meaningful policy interventions. A "word gap" in early childhood language exposure, for example, suggests much clearer interventions than a persistent "black–white test score gap." More broadly, the challenges posed by ethnic conflict and racial inequality are much more likely to be understood and addressed if scholars disaggregate the elements of race and identify the particular ways difference is turned into disparity.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

We are grateful to Bear Braumoeller, Jennifer Bréa, John Bullock, Kevin Clarke, Adam Glynn, Jim Greiner, Jennifer Hochschild, Luke Keele, Gary King, Rich Nielsen, Kevin Quinn, Eric Schickler, Shauna Shames, Ali Valenzuela, and Teppei Yamamoto for thoughtful advice and suggestions. We also thank participants at the 2011 Midwest Political Science Association panel on causal inference for helpful feedback. Authors' names are in alphabetical order.

LITERATURE CITED

Abdelal R, Herrera YM, Johnston AI, McDermott R. 2009. Identity as a variable. In *Measuring Identity*, ed. R Abdelal, YM Herrera, AI Johnston, R McDermott, pp. 17–32. New York: Cambridge Univ. Press

Abrajano M, Panagopoulos C. 2011. Does language matter? The impact of Spanish versus English-language GOTV efforts on Latino turnout. *Am. Polit. Res.* 39(4):643–63

Adida CL, Laitin DD, Valfort M-A. 2010. Identifying barriers to Muslim integration in France. *PNAS* 107(52):1–7

Almond D, Hoynes HW, Schanzenbach DW. 2011. Inside the war on poverty: the impact of food stamps on birth outcomes. *Rev. Econ. Stat.* 93(2):387–403

Angrist JD, Imbens GW, Rubin DB. 1996. Identification of causal effects using instrumental variables. *J. Am. Stat. Assoc.* 91(434):444–55

Angrist JD, Pischke J-S. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ: Princeton Univ. Press

Appiah KA. 1985. The uncompleted argument: Du Bois and the illusion of race. *Crit. Inquiry* 12(1):21–37

Bar-Haim Y, Ziv T, Lamy D, Hodes RM. 2006. Nature and nurture in own-race face processing. *Psychol. Sci.* 17(2):159–63

Bedolla LG, Michelson MR. 2012. *Mobilizing Inclusion: Transforming the Electorate Through Get-Out-the-Vote Campaigns*. New Haven, CT: Yale Univ. Press

Behrman JR, Rosenzweig MR. 2004. Returns to birthweight. *Rev. Econ. Stat.* 86(2):586–601

Ben-Shlomo Y, Kuh D. 2002. A life course approach to chronic disease epidemiology: conceptual models, empirical challenges and interdisciplinary perspectives. *Int. J. Epidemiol.* 31(2):285–93

Bertrand M, Mullainathan S. 2004. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *Am. Econ. Rev.* 94(4):991–1013

Blank RM, Dabady M, Citro CF, eds. 2004. *Measuring Racial Discrimination*. Washington, DC: Natl. Acad. Press

Bobo LD, Johnson D. 2004. A taste for punishment: black and white Americans' views on the death penalty and the war on drugs. *Du Bois Rev. Soc. Sci. Res. Race* 1(1):151–80

Boker SM, Cohn JF, Theobald B-J, Matthews I, Mangini M, et al. 2011. Something in the way we move: motion dynamics, not perceived sex, influence head movements in conversation. *J. Exp. Psychol. Hum. Percept. Perform.* 37(3):874–91

Broockman DE. 2013. Black politicians are more intrinsically motivated to advance blacks' interests: a field experiment manipulating political incentives. *Am. J. Polit. Sci.* 57(3):521–36

Butler DM, Broockman DE. 2011. Do politicians racially discriminate against constituents? A field experiment on state legislators. *Am. J. Polit. Sci.* 55(3):463–77

Chandra K. 2006. What is ethnic identity and does it matter? *Annu. Rev. Polit. Sci.* 9(1):397–424

Conley D, Strully KW. 2012. Birth weight, infant mortality, and race: twin comparisons and genetic/environmental inputs. *Soc. Sci. Med.* 75(12):2446–54

Cook TD, Campbell DT. 1979. *Quasi-Experimentation: Design & Analysis Issues for Field Settings*. Boston: Houghton Mifflin
Crenshaw K. 1991. Mapping the margins: intersectionality, identity politics, and violence against women of color. *Stanford Law Rev.* 43(6):1241–99
Cutler DM, Fryer RG, Glaeser EL. 2005. *Racial differences in life expectancy: the impact of salt, slavery, and selection*. March 1. **http://isites.harvard.edu/fs/docs/icb.topic98848.files/salt_science_submission_3-01.pdf**
Dawson MC. 1994. *Behind the Mule: Race and Class in African-American Politics*. Princeton, NJ: Princeton Univ. Press
Deaux K. 1985. Sex and gender. *Annu. Rev. Psychol.* 36(1):49–81
Dehejia RH, Wahba S. 2002. Propensity score-matching methods for nonexperimental causal studies. *Rev. Econ. Stat.* 84(1):151–61
Devlin B. 1997. *Intelligence, Genes and Success: Is It All in the Genes?: Scientists Respond to the Bell Curve*. New York: Springer
Dolnick S. 2011. Ethnic differences emerge in plastic surgery. *N. Y. Times*, Feb. 18. **http://www.nytimes.com/2011/02/19/nyregion/19plastic.html**
Du Bois WEB. 2007 (1903). *The Souls of Black Folk*. New York: Oxford Univ. Press
Duster T. 2005. Race and reification in science. *Science* 307(5712):1050–51
Enos RD. 2011. *Reconsidering racial threat: experimenting on the contextual causes of intergroup attitudes*. Presented at Lauder School of Government, Diplomacy and Strategy, Program in Decision Making and Political Psychology, Harvard–IDC Symp. Political Psychology and Decision Making, IDC Herzliya, Israel, Nov.
Faulkner W. 1990 (1932). *Light in August*. London: Vintage
Freeman JB, Penner AM, Saperstein A, Scheutz M, Ambady N. 2011. Looking the part: social status cues shape race perception. *PLoS ONE* 6(9):e25107
Gates HL Jr. 1997. The passing of Anatole Broyard. In *Thirteen Ways of Looking at a Black Man*, pp. 180–214. New York: Vintage
Gay C. 2012. Moving to opportunity: the political effects of a housing mobility experiment. *Urban Aff. Rev.* 48(2):147–79
Gay C, Hochschild JL. 2010. *Is racial linked fate unique? Comparing race, ethnicity, class, gender, and religion*. Presented at Annu. Meet. Am. Polit. Sci. Assoc., Washington, DC, Sep. 4
Gelman A, Hill J. 2007. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. New York: Cambridge Univ. Press
Gilens M. 1996. "Race coding" and white opposition to welfare. *Am. Polit. Sci. Rev.* 90(3):593–604
Gilliam FD Jr, Iyengar S. 2000. Prime suspects: the influence of local television news on the viewing public. *Am. J. Polit. Sci.* 44(3):560–73
Green DP. 2004. Mobilizing African-American voters using direct mail and commercial phone banks: a field experiment. *Polit. Res. Q.* 57(2):245–55
Greenwald AG, McGhee DE, Schwartz JLK. 1998. Measuring individual differences in implicit cognition: the Implicit Association Test. *J. Pers. Soc. Psychol.* 74(6):1464–80
Greiner DJ, Rubin DB. 2010. Causal effects of perceived immutable characteristics. *Rev. Econ. Stat.* 93(3):775–85
Griffin JH. 1996 (1962). *Black Like Me*. New York: NAL Trade
Guttentag M, Secord PF. 1983. *Too Many Women?: The Sex Ratio Question*. Thousand Oaks, CA: Sage
Halsell G. 1969. *Soul Sister*. New York: World
Haslam N, Rothschild L, Ernst D. 2000. Essentialist beliefs about social categories. *Br. J. Soc. Psychol.* 39(1):113–27
Heckman JJ. 1998. Detecting discrimination. *J. Econ. Perspect.* 12(2):101–16
Heckman JJ, Siegelman P. 1993. The urban institute audit studies: their methods and findings. In *Clear and Convincing Evidence: Measurement of Discrimination in America*, ed. M Fix, RJ Struyk, pp. 187–258. Washington, DC: Urban Inst.
Hernstein R, Murray C. 1994. *The Bell Curve: Intelligence and Class Structure in American Life*. New York: Simon & Schuster
Hochschild JL, Powell BM. 2008. Racial reorganization and the United States Census 1850–1930: mulattoes, half-breeds, mixed parentage, Hindoos, and the Mexican race. *Stud. Am. Polit. Dev.* 22(1):59–96

Hochschild JL, Sen M. 2015. Technology optimism or pessimism about genomic science: variation among experts and scholarly disciplines. *Ann. Am. Acad. Polit. Soc. Sci.* 658(1):236–52

Hochschild JL, Weaver V. 2010. "There's no one as Irish as Barack O'Bama": the politics and policy of multiracialism in the United States. *Perspect. Polit.* 8(3):737–59

Holland PW. 1986. Statistics and causal inference. *J. Am. Stat. Assoc.* 81(396):945–60

Holland PW. 2008. Causation and race. In *White Logic, White Methods: Racism and Methodology*, ed. T Zuberi, E Bonilla-Silva, pp. 93–109. Lanham, MD: Rowman & Littlefield

Hopkins D. 2015. The upside of accents: language, skin tone, and attitudes toward immigration. *Br. J. Polit. Sci.* 45(3):531–57

Htun M. 2005. What it means to study gender and the state. *Polit. Gender* 1(01):157–66

Huber GA, Lapinski JS. 2006. The "race card" revisited: assessing racial priming in policy contests. *Am. J. Polit. Sci.* 50(2):421–440

Imai K, Keele L, Tingley D, Yamamoto T. 2011. Unpacking the black box of causality: learning about causal mechanisms from experimental and observational studies. *Am. Polit. Sci. Rev.* 105(04):765–89

Imai K, King G, Stuart EA. 2008. Misunderstandings between experimentalists and observationalists about causal inference. *J. R. Stat. Soc. Ser.* 171(2):481–502

Imbens GW, Rubin DB. 2010. *Causal Inference in Statistics and Social Sciences*. Draft book manuscript, Dep. Econ., Dep. Statistics, Harvard Univ.

James M. 2011. Race. In *The Stanford Encyclopedia of Philosophy*, ed. EN Zalta. **http://plato.stanford.edu/entries/race/**

Junn J, Masuoka N. 2008. Asian American identity: shared racial status and political context. *Perspect. Polit.* 6(4):729–40

Katz LF, Kling JR, Liebman JB. 2001. Moving to opportunity in Boston: early results of a randomized mobility experiment. *Q. J. Econ.* 116(2):607–54

Keele L. 2010. *An overview of rbounds: an R package for Rosenbaum bounds sensitivity analysis with matched data*. White Pap., Dep. Polit. Sci., Penn. State Univ, Columbus, OH

Kelly DJ, Quinn PC, Slater AM, Lee K, Gibson A, et al. 2005. Three-month-olds, but not newborns, prefer own-race faces. *Dev. Sci.* 8(6):F31–F36

Kennedy R. 2012. *Interracial Intimacies: Sex, Marriage, Identity, and Adoption*. New York: Random House

Kim CJ, Lee T. 2001. Interracial politics: Asian Americans and other communities of color. *PS: Polit. Sci. Polit.* 34(3):631–37

King G. 1991. "Truth" is stranger than prediction, more questionable than causal inference. *Am. J. Polit. Sci.* 35(4):1047–53

King G, Keohane RO, Verba S. 1994. *Designing Social Inquiry: Scientific Inference in Qualitative Research*. Princeton, NJ: Princeton Univ. Press

King G, Zeng L. 2006. The dangers of extreme counterfactuals. *Polit. Anal.* 14(2):131–51

Kitcher P. 2007. Does 'race' have a future? *Philos. Public Aff.* 35(4):293–317

Knowles ED, Peng K. 2005. White selves: conceptualizing and measuring a dominant-group identity. *J. Pers. Soc. Psychol.* 89(2):223–41

Kurzban R, Tooby J, Cosmides L. 2001. Can race be erased? Coalitional computation and social categorization. *PNAS* 98(26):15387–92

Kuzawa CW, Sweet E. 2009. Epigenetics and the embodiment of race: developmental origins of US racial disparities in cardiovascular health. *Am. J. Hum. Biol.* 21(1):2–15

Lee T. 2008. Race, immigration, and the identity-to-politics link. *Annu. Rev. Polit. Sci.* 11:457–78

Lee T, Pérez EO. 2014. The persistent connection between language-of-interview and Latino political opinion. *Polit. Behav.* 36(2):401–25

Liebler CA, Rastogi S, Fernandez LE, Noon JM, Ennis SR. 2014. *America's churning races: race and ethnic response changes between Census 2000 and the 2010 Census*. Cent. Admin. Rec. Res. Applic. Work. Pap. Ser. #2014–09. **https://www.census.gov/srd/carra/Americas_Churning_Races.pdf**

López IFH. 1994. The social construction of race: some observations on illusion, fabrication, and choice. *Harvard Civil Rights–Civil Liberties Law Rev.* 29(1):1–62

Loury GC. 2002. *The Anatomy of Racial Inequality*. Cambridge, MA: Harvard Univ. Press

Mendelberg T. 2001. *The Race Card: Campaign Strategy, Implicit Messages, and the Norm of Equality*. Princeton, NJ: Princeton Univ. Press
Miller JM, Krosnick JA. 2000. News media impact on the ingredients of presidential evaluations: politically knowledgeable citizens are guided by a trusted source. *Am. J. Polit. Sci.* 44(2):301–15
Morning A. 2011. *The Nature of Race: How Scientists Think and Teach About Human Difference*. Berkeley: Univ. Calif. Press
Navara KJ. 2009. Humans at tropical latitudes produce more females. *Biol. Lett.* 5(4):524
Nisbett RE, Cohen D. 1996. *Culture of Honor: The Psychology of Violence in the South*. Boulder, CO: Westview
Nobles M. 2000. *Shades of Citizenship: Race and the Census in Modern Politics*. Stanford, CA: Stanford Univ. Press
Noymer A, Penner AM, Saperstein A. 2011. Cause of death affects racial classification on death certificates. *PLOS ONE* 6(1):e15812
O'Connor M. 2014. Is race plastic? My trip into the "ethnic plastic surgery" minefield. *N. Y. Mag.*, July 27. **http://nymag.com/thecut/2014/07/ethnic-plastic-surgery.html**
Omi M, Winant H. 1994. *Racial Formation in the United States: From the 1960s to the 1990s*. New York: Routledge
Pager D. 2003. The mark of a criminal record. *Am. J. Sociol.* 108(5):937–75
Pager D. 2007. The use of field experiments for studies of employment discrimination: contributions, critiques, and directions for the future. *Ann. Am. Acad. Polit. Soc. Sci.* 609(1):104–33
Pearl J. 2000. *Causality: Models, Reasoning, and Inference*. New York: Cambridge Univ. Press
Penner AM, Saperstein A. 2008. How social status shapes race. *PNAS* 105(50):19628–30
Purnell T, Idsardi W, Baugh J. 1999. Perceptual and phonetic experiments on American English dialect identification. *J. Lang. Soc. Psychol.* 18(1):10–30
Rosenbaum PR. 2002. *Observational Studies*. New York: Springer Verlag
Rubin DB. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *J. Educ. Psychol.* 66(5):688–701
Rubin DB. 1978. Bayesian inference for causal effects: the role of randomization. *Ann. Stat.* 6(1):34–58
Rubin DB. 2005. Causal inference using potential outcomes. *J. Am. Stat. Assoc.* 100(469):322–31
Rutter M, Tienda M. 2005. The multiple facets of ethnicity. In *Ethnicity and Causal Mechanisms*, ed. M Rutter, M Tienda, pp. 50–79. New York: Cambridge Univ. Press
Saperstein A. 2006. Double-checking the race box: examining inconsistency between survey measures of observed and self-reported race. *Soc. Forces* 85(1):57–74
Saperstein A, Penner AM. 2010. The race of a criminal record: how incarceration colors racial perceptions. *Soc. Probl.* 57(1):92–113
Scarr S, Pakstis AJ, Katz SH, Barker WB. 1977. Absence of a relationship between degree of white ancestry and intellectual skills within a black population. *Hum. Genet.* 39(1):69–86
Schuyler GS. 1971 (1931). *Black No More*. New York: Dover
Sekhon JS. 2009. Opiates for the matches: matching methods for causal inference. *Annu. Rev. Polit. Sci.* 12:487–508
Sharkey P. 2010. The acute effect of local homicides on children's cognitive performance. *PNAS* 107(26):11733–38
Sidanius J, Pratto F. 2001. *Social Dominance: An Intergroup Theory of Social Hierarchy and Oppression*. New York: Cambridge Univ. Press
Sniderman PM, Piazza TL. 1993. *The Scar of Race*. Cambridge, MA: Belknap
Splawa-Neyman J, Dabrowska DM, Speed TP. 1990. On the application of probability theory to agricultural experiments. Essay on principles. Sec. 9. *Stat. Sci.* 5(4):465–72
Steele C. 1997. A threat in the air: how stereotypes shape intellectual identity and performance. *Am. Psychol.* 52(6):613–29
Tesler M. 2012. The spillover of racialization into health care: how President Obama polarized public opinion by racial attitudes and race. *Am. J. Polit. Sci.* 56(3):690–704
Thomas K, Zarda B. 2010. In N.C.A.A., question of bias over a test for a genetic trait. *N. Y. Times*, Apr. 12. **http://www.nytimes.com/2010/04/12/sports/12sickle.html**

Valentino NA, Hutchings VL, White IK. 2002. Cues that matter: how political ads prime racial attitudes during campaigns. *Am. Polit. Sci. Rev.* 96(01):75–90

Valenzuela AA, Michelson MR. 2011. *Turnout, status and identity: mobilizing Latinos to vote in contrasting contexts*. Presented at Annu. Meet. Western Polit. Sci. Assoc., San Antonio, TX, Apr. 21–23

VanderWeele TJ, Hernán MA. 2012. Causal effects and natural laws: towards a conceptualization of causal counterfactuals for nonmanipulable exposures, with application to the effects of race and sex. In *Causality: Statistical Perspectives and Applications*, pp. 101–13. Hoboken, NJ: Wiley

Walton GM, Cohen GL. 2011. A brief social-belonging intervention improves academic and health outcomes of minority students. *Science* 331(6023):1447–51

Wasow O. 2012. *Violence and voting: Did the 1960s black uprisings reshape American politics?* Presented at Annu. Meet. Midwest Polit. Sci. Assoc., Chicago, IL, Apr. 12–15

West C, Zimmerman DH. 1987. Doing gender. *Gender Soc.* 1(2):125–51

White I. 2007. When race matters and when it doesn't: racial group differences in response to racial cues. *Am. Polit. Sci. Rev.* 101(2):339–54

Winship C, Morgan SL. 1999. The estimation of causal effects from observational data. *Annu. Rev. Sociol.* 25:659–706

Zuberi T. 2001. *Thicker Than Blood: How Racial Statistics Lie*. Minneapolis: Univ. Minn. Press

Zuberi T, Bonilla-Silva E. 2008. *White Logic, White Methods: Racism and Methodology*. Lanham, MD: Rowman & Littlefield

Zuckerman M. 1990. Some dubious premises in research and theory on racial differences: scientific, social, and ethical issues. *Am. Psychol.* 45(12):1297–303

Annual Review of Political Science
Volume 19, 2016

Contents

Indexes

Errata

An online log of corrections to *Annual Review of Political Science* articles may be found at http://www.annualreviews.org/errata/polisci