Journal of Social Issues, Vol. 68, No. 2, 2012, pp. 385–398

Intergroup Disparities and Implicit Bias: A Commentary

Emily L. Fisher* **and Eugene Borgida**
University of Minnesota

Despite changing societal norms that are less tolerant of overt discrimination, demonstrable disparities between racial and gender groups remain. The contributors to this impressive special issue document and offer explanations for these disparities in employment and professional contexts, and with regard to disparate outcomes associated with the legal system, and in medical and health care contexts. In our commentary, we examine these aggregate-level disparities and the individual-level explanatory accounts proposed for their existence. The evidence that these papers present is often enough to rule out alternative explanatory accounts, and implicit bias remains a viable account for disparities that, to varying degrees, fit with the available data and the claim that implicit biases are contributing to an understanding of ongoing real-world disparities. As such, we believe that implicit bias research will continue to play a crucial role in understanding and hopefully reducing these aggregate-level disparities in employment, legal, and health care domains.

Despite changing societal norms that are less tolerant of overt discrimination, disparities between racial and gender groups remain. The authors in this special issue document these disparities and explain how and why they have been perpetuated. Contemporary discrimination persists in a range of domains that have very real and serious consequences for members of disadvantaged groups; the focus on employment, legal, and health care disparities provides some exemplars of an insidious problem. In our commentary, we examine these aggregate-level disparities and the explanatory accounts for their existence proposed by the distinguished set of contributors to this special issue. We ask: which accounts have been conclusively dismissed, and which remain viable contenders for the cause of group-based disparities? We pursue this discussion in the context of the three

*Correspondence regarding this article should be addressed to Emily L. Fisher, Department of Psychology, Hobart & William Smith Colleges, Geneva, NY 14456; e-mail: fisher@hws.edu.

domains that characterize the papers included in the special issue: disparities in employment and professional contexts, disparate outcomes associated with the legal system, and medical and health care disparities.

Disparities in Professional and Employment Contexts

Cikara, Rudman, and Fiske (2012) examine gender differences in academics' publication rates in the *Journal of Personality and Social Psychology* (*JPSP*). They suggest that men publish in *JPSP* for a longer portion of their careers and do so more frequently, and the investigators seek to determine what factors could account for this gap. A variety of factors are examined, including gender bias, differences in research quality, effort or persistence, family demands, and institutional status of the author's PhD-granting institution as well as current place of employment. However, the data they collect do not support most of these explanatory accounts. When there are gender differences in these variables, the benefit(s) often went to the man. For instance, when examining authors' persistence, Cikara et al. note that negotiating with editors led to more publications for men, but there was no such effect for women. Similarly, marriage and family demands were actually associated with more *JPSP* articles for men, but had no effect for women. The women in this sample had received their PhDs from (on average) more prestigious institutions than the men, so this variable cannot account for the publication gap. After each of the accounts are examined and dismissed, the one remaining explanatory account is a subtle gender bias within the field. Cikara et al. are careful to note that they do not have the kind of direct evidence for this bias needed to establish this account for the disparity, but when other potential accounts are dismissed, it is hard to avoid attributing this gap to anything else.

The other two groups of contributors to the employment contexts section of this special issue concern themselves with employment disparities that affect a wider range of career types: they focus on racial disparities in hiring job applicants. Pager and Western (2012) provide detailed documentation of the existence of such disparities. The authors present results of two field studies which used a technique called "matched pair testing" in which pairs of testers with different racial status but equivalent qualifications apply for the same actual job openings. With such experimental controls to rule out any confounding factors that could affect employment attainment, they were able to demonstrate that employers appear to use different evaluation standards for White and minority applicants. Case studies of some of their tester pairs indicate that employers are not categorically denying opportunities to minorities; rather, they simply seem to allow a bit more flexibility in their evaluation of White applicants and maintain stricter standards in the evaluation of minority applicants. These data clearly establish that discrimination (in the form of fewer interviews and job offers for minorities compared to equally

qualified White candidates) occurs at the aggregate level, and imply that racial bias must be an explanatory factor to be taken seriously. However, this type of data does not allow one to draw conclusions about any particular employer's reasons for discriminating, since "even a non-discriminatory employer, when forced to choose between two equally qualified candidates, will choose the white applicant in half of all cases" (Pager & Western, 2012, p. 233). Once again, it is impossible to rule out implicit bias as an account for a demonstrable disparity, but the data do not yet allow us to rule it in with causal certainty.

Bendick and Nunes (2012) concur with Pager and Western that hiring bias exists, and also espouse the power of matched pair testing for documenting these disparities without the need to generalize from laboratory results. They review the psychological literature supporting two potential accounts for this disparity. First, stereotypes about a racial outgroup (whether explicit or implicit) can automatically affect impression formation, memory, and other processes that influence judgments of people. Second, when an applicant is a member of the employer's own racial group, ingroup preferences can bias judgments in favor of the applicant. Bendick and Nunes note that the hiring process can be especially vulnerable to these biases because of the characteristics of the situation: there is limited individuating information available about a job applicant, decisions are often made under time pressure, and (compared to other employment decisions such as promotions) there is less of a chance to reexamine a decision and correct for errors. Because race and gender hiring and pay disparities exist even when controlling for a myriad of factors that could influence employers' decisions, automatic stereotyping and ingroup favoratism cannot be ruled out as accounts for such disparities. Bendick and Nunes praise and advocate for the matched pair testing technique insofar as it provides powerful evidence that the disparity needs to be addressed. But what this methodology ultimately cannot overcome is the need to infer what is happening inside employers' minds: it cannot inform us of the psychological processes underlying the hiring bias.

In sum, when evaluating the conclusions that all of these authors reach, one common theme clearly emerges. Implicit bias cannot be definitively ruled out as a cause of these group-based disparities. When researchers examine other accounts for the disparities discussed throughout this special issue, they can provide evidence demonstrating that each of these explanations cannot scientifically account for the gaps. However, it is almost a truism in science that we can disprove hypotheses, but we can never really prove them. Thus, as good scientists, none of the contributors to this issue leap to conclusions from their data to claim that implicit bias must be the cause of the disparity. Many strongly suspect that it is, but are impressively circumspect about the extent to which their data allow them to definitively rule it in for all situations. Ruling out *and* ruling in explanations can both help us to tighten and strengthen the causal reasoning about the role of implicit bias in these disparities.

Disparities in the Legal Context

A second set of disparities examined in this special issue involves the legal system, and the experience of racial groups in various legal contexts. For example, Sadler, Correll, Park, and Judd (2012) point out that minorities are incarcerated and shot at by police at rates disproportionate to their numbers in the general population. It may be that minorities simply commit a disproportionate number of crimes, and Sadler et al. admit that evidence regarding the veracity of this claim is mixed. However, it may also be that police officers' behavior is affected by the stereotype that associates minorities with criminal activity. To test this hypothesis, they review recent theoretical developments that rely on the "shooter bias" paradigm (Correll, Park, Judd, & Wittenbrink, 2002; Correll et al., 2007) and present three new studies that extend this methodology. The extant literature demonstrates that police tend to associate Blacks with weapons, based on their reaction times and accuracy on the shooter bias task. Sadler et al. (2012) teach us that this "shooter bias" extends to Latino targets as well, and that police officers' explicit endorsement of stereotypes about racial tendencies to commit crimes relates to the officers' performance on the shooter bias task. These new data reinforce the idea that police officers may have an automatic association with violence for some racial minorities, so we cannot rule out this bias as an account for the disparities in the tendency to suspect members of certain races that go beyond actual base rates of criminality. However, as Sadler et al. point out, their data cannot empirically establish the reasons for this association. Perhaps the association could be a reflection of the officer's experiences on the job, or it could be factoring into self-fulfilling prophecies during officers' interactions with minorities (i.e., officers expect a minority suspect to be more violent and therefore unintentionally treat the suspect in a way that provokes him or her to then react more violently). Thus, more research is needed before we claim that the "shooter bias" is due to police officers' prejudice, although this account cannot be dismissed.

Disparities in Health Care Contexts

In this section of the special issue on health care contexts, Eggly, Griggs, Orom, Penner, and Underwood (2012) document several alarming racial disparities in medical treatments: Black patients receive less information from their oncology doctors than White patients do, Black breast cancer patients are less likely to receive the proper doses for chemotherapy and other drug regimens, and Black prostate cancer patients are less likely to undergo surgery or other definitive treatments for their cancer. Each of these disparities contributes to differential survival and mortality rates for Black and White patients with cancer. Eggly et al. examine several potential explanations for these disparities. At the systemic level, racial disparities could be due to broader gaps in socioeconomic status and

access to public health services that correlate with race. If Blacks are less likely to be insured, and tend to go to poorer quality hospitals, then one would expect poorer outcomes. Disparities could also be due to patient factors: Black patients might be less likely to trust their doctors and thus be more likely to be medically noncompliant. Finally, there could be provider-level factors at play. The physicians themselves could fall prey to subtle prejudices and stereotypes that lead them to treat patients differently based on race. However, evidence discounts the first of these accounts. When Eggly et al. examine the issue at the systemic level, they find that disparities cannot just be due to socioeconomic differences because race-based disparities persist even when controlling for such factors. For instance, even when Blacks and Whites are on the same insurance plan, their health outcomes are not equal. Although there is evidence that patients' trust of their doctors and perceptions of discrimination do predict the degree to which they adhere to treatment plans and seek preventive care, Eggly et al. caution against blaming the victim in this case. After all, Black patients' attitudes may very well reflect an understandable reaction to the institutional and everyday racism that they face in society. So even if patient attitudes are part of the explanation for health disparities, they cannot be the ultimate cause. This leaves us with physicians' attitudes as an account for health disparities, and evidence does suggest that physician bias may well play a role. Several recent studies (Green et al., 2007; Penner et al., 2010) suggest that physicians do hold implicit biases that affect their treatment recommendations and that patients respond negatively to physicians who can be classified as aversive racists (i.e., they have negative implicit attitudes despite explicit attitudes that are egalitarian). Eggly et al. also highlight differences in doctor–patient interactions. For Blacks, these interactions are much more likely to be interracial interactions than they are for Whites, and research in medical settings (Penner, Albrecht, Orom, Coleman, & Underwood, 2009) suggests that such interactions are shorter and less positive. In sum, Eggly et al. do not rule out physician bias (especially as it leads to doctor–patient communication problems as an account for racial disparities in cancer treatments).

Another disparity in the health context is highlighted in this issue: racial disparities in social capital and well being. For example, Brondolo, Libretti, Riviera, and Walsemann (2012) review findings at different levels of analysis which suggest that racism adversely affects the development and durability of cross-race peer relationships (a key component, they argue, of social capital). In particular, Brondolo et al. discuss the ways in which racism could be a psychological stressor that constrains and limits the social capital of minorities, which in turn makes them more vulnerable to disease (as well as to other disparate economic and social outcomes) relative to members of groups that do not face racism regularly. Although they maintain that many other factors contribute to social and physical well being, they review data that support a connection between racism indicators of social capital, and cardiac health. For instance, there is a relationship between

experiencing racism and risk factors for cardiac disease such as depression and smoking. Brondolo et al. are not yet ready to claim that racism is an indisputable factor in cardiac health (i.e., racism may mediate the direct effects of racism on health outcomes), but they do suggest that it is an important area for further research. Implicit bias as an explanatory account for such health care disparities, or as an account for disparities in the employment context, has generated considerable scientific scrutiny. Some of these issues and concerns, discussed briefly in the next section, focus on the use of the Implicit Association Test (IAT) as a measure of implicit bias.

Implicit Bias and Intergroup Disparities

Despite (or perhaps because of) the novel and important claims that implicit bias research can make, the study of implicit bias has attracted its fair share of controversy, especially with regard to measures of implicit bias. As this special issue demonstrates, implicit bias is a potential factor in many important contexts, and its effects have been examined from a variety of methodologies and perspectives. Nevertheless, critics sometimes conflate certain measures and methods (most notably, the Implicit Association Test or IAT; Greenwald, McGhee, & Schwartz, 1998) with the scientific study of implicit bias, more generally. We briefly turn our focus to two recent examples that highlight this controversy and point in the direction that research needs to go in order to most effectively understand implicit bias in non-laboratory contexts. Both are cases in which psychological journals have published commentaries from critics of the IAT and implicit bias research along with responses from proponents of such research. These exchanges have highlighted the concerns that implicit bias research raises as well as attempts to address such concerns. Studying this debate can illuminate some relevant concerns for this issue's theme of the realities of contemporary discrimination and reveals the degree to which social scientists have reached consensus regarding implicit bias as a scientific account for discriminatory actions.

Major Criticisms of Implicit Bias Research

In a recent issue of *Research in Organizational Behavior*, Tetlock and Mitchell's (2009) latest critique of the IAT and implicit bias research (see also Arkes & Tetlock, 2004; Mitchell & Tetlock, 2006; Tetlock & Arkes, 2004; Tetlock, Mitchell, & Murray 2008) focuses on three broad criticisms of such research programs. First, they question the construct validity of the IAT and propose that other factors, such as familiarity with the stimuli, anxiety over the test, sympathy for disadvantaged groups, or cognitive dexterity, may actually be what the IAT is measuring (rather than levels of prejudice, per se). They also suggest that

the IAT's test-retest reliability and its varying correlations with explicit measures of attitudes indicate a high degree of error variance among individuals' scores.

Another point of contention for Tetlock and Mitchell (2009) is the psychometric issue of whether a person with a given IAT score will be more or less likely to discriminate. This criticism is based on arguments such as Blanton and Jaccard's (2006) discussion of "arbitrary metrics": it is unknown, they argue, how a given reaction time will map out onto observable behavior. They also propose that the size of the IAT's correlations with behavioral measures will lead to incredibly high rates of false positives (i.e., a nonprejudiced person being accused of prejudice based on "failing" the IAT).

Finally, they take on the question of external validity, and argue that implicit bias research will not generalize to real workplaces given the differences between such situations and the typical psychology laboratory. This point of view is shared by other IAT skeptics; for instance, Landy's (2008) target article in *Industrial and Organizational Psychology: Perspectives on Science and Practice* focuses on the application of laboratory-based stereotyping research to real-life employment decision making, and concludes that it is often inappropriate to make inferential leaps between these contexts. He outlines numerous differences between the typical lab study and an employment setting, mainly focused around the lack or presence of individuating information, and the degree of accountability (Lerner & Tetlock, 1999). Because individuating information has been shown, theoretically and empirically, to decrease the tendency to rely on stereotypes (e.g., Fiske & Neuberg, 1990; Kunda & Sinclair, 1999), and people in the real world tend to have individuating information about their subordinates, Landy (2008) argues (but does not establish empirically) that stereotyping will have only a small or insignificant effect on real-world decision making in complex organizations.

Specific Attempts to Address These Criticisms

Such criticisms have not gone unanswered (e.g., Borgida, Deason, Kim, & Fiske, 2008; Jost et al., 2009; Kang, 2012). Some of this debate has actually stimulated scientific work to clarify and improve the status of the IAT (Greenwald, Poehlman, Uhlmann, & Banaji, 2009; Nosek, Greenwald & Banaji, 2005, 2007). Moreover, many psychologists accept the basic claims of implicit bias or implicit prejudice and are concerned about their implications for society. For instance, Jost et al. (2009) paint a picture that puts Tetlock and Mitchell (2009) in the role of the global warming skeptics of the implicit bias world—a minority railing against sound scientific consensus. Jost et al. (2009) make the case that the IAT is just one methodology to emerge from decades' worth of empirical evidence and theoretical development in social and cognitive psychology; one that has evolved in an evidence-based way. Questioning the explanatory and predictive power of implicit racial attitudes, they argue, would be tantamount to dismissing

much of modern psychology: much established knowledge about social cognition would need to be similarly disregarded on the basis of Tetlock and Mitchell's (2009) stance. Moreover, even in the applied setting of the courtroom, scientific testimony about implicit bias and stereotyping is routinely accepted (Borgida et al., 2008).

Several commentaries provide support for the IAT's validity and reliability in a variety of contexts. Rudman (2008; see also Greenwald, Poehlman, Uhlmann, & Banaji, 2009; Nosek, Greenwald, & Banaji, 2007) provides an overview of the IAT's validity as demonstrated by a wealth of research. Although early versions of the IAT were somewhat more influenced by constructs unrelated to bias such as cognitive skill, its scoring procedure has been reevaluated and updated to make it more robust to extraneous variance and more validly associated with implicit emotional conditioning (Nosek et al., 2007). The IAT's predictive validity has also been demonstrated in a wide range of contexts with many real-world behaviors, such as doctors' treatment recommendations for patients (Eggly et al., 2012; Green et al., 2007), shooter biases (Glaser & Knowles, 2008; Sadler et al., 2012), evaluations of racial minority targets (Amodio & Devine, 2006; Rudman & Lee, 2002), and actual hiring decisions (Rooth, 2010). In contrast to Landy's (2008) critique that the IAT only predicts micro-level behaviors in laboratory settings, work on the IAT's validity has expanded to include the prediction of more macro-behaviors in real-world samples (Greenwald, 2008; Jost et al., 2009). Other researchers have worked to establish the IAT's construct validity through neurological and psychophysiological correlates. For instance, IAT scores are associated with amygdala activation when exposed to other-race photos (Phelps et al., 2000) and with physiological stress markers during interracial encounters (Mendes, Gray, Mendoza-Denton, Major, & Epel, 2007).

Despite criticisms that the IAT is not relevant outside of the laboratory (Landy, 2008) many studies that use the IAT or other measures of implicit bias are conducted in such applied settings, providing corroboration for the implicit bias account using real-world behaviors and a variety of samples. For instance, Jost et al. (2009; see also Greenwald, 2008) highlight 10 recent studies that administered IATs to diverse samples of professionals, from doctors and nurses to police officers and employment recruiters. In each of these cases, IAT scores predicted a potentially discriminatory behavioral outcome. For example, Rooth (2010) asked employment recruiters to complete an IAT. An audit study sending resumes to these same recruiters (in response to actual job openings in a Swedish city) revealed an association between implicit bias and a preference to call back applicants with a Swedish surname over a Turkish surname.

Regardless of any debate over IAT validity, the broader point that is often lost amidst the methodological and ideological cacophony is that considerable implicit bias research goes beyond the IAT and uses methods that have been regarded with less criticism. The IAT, after all, is only one of several families of implicit bias measures that have been adapted from social cognition research more

broadly; others include affective priming or linguistic analysis (Fazio & Olson, 2003). Research using these types of techniques also finds that implicit bias predicts a variety of behavioral outcomes in intergroup domains. The "shooter bias" paradigm (Correll et al., 2002; Sadler et al., 2012) offers a method that relies on reactions to ambiguously threatening targets and demonstrates the extent to which people's judgments are in line with racial stereotypes. Other research paradigms move discrimination research entirely outside of the laboratory and into the field. As several of the special issue authors discuss, audit studies are a prime example of this research (Bendick & Nunes, 2012; Pager & Westen, 2012; see also Leslie, King, Bradley, & Hebl, 2008; Pager, 2003; Pager & Quillian, 2005). Researchers can send resumes or confederates to apply for open positions. This strategy allows for the applicants to have equivalent qualifications for the job, yet vary on group membership. When the applicant with one group membership receives more interviews or job offers than the applicant with another group membership, managers' bias against that group is more likely to be the only explanation for the discrepancy. Similarly, work outcome studies use employers' data on salaries, promotions, or other outcomes, and perform multivariate analyses to control for legitimate causes of variance (Leslie et al., 2008). When disparities between groups remain despite such controls, the results are regarded as evidence of discrimination or bias in employment decisions and enhance our ability to rule in or rule out implicit bias as a causal account for the disparities in question.

Another way to address criticisms of IAT and implicit bias research is by challenging the critics' assumptions about the context of discrimination. For instance, Landy (2008) believes that such research is overblown because stereotyping effects often diminish in the presence of individuating information about a person in complex work contexts. Thus, in a real-life employment situation, such individuating information is likely to take precedence over categorical stereotypes. However, this perspective ignores the many obstacles to using such individuating information in the real world. As Bendick & Nunes (2012) demonstrate, hiring decisions are frequently made in an environment in which very little individuating information is available: a résumé or brief interview may not provide much information about a job candidate. Moreover, if that applicant's group membership activates any stereotypes for the hiring decision maker, those stereotypes are likely to affect the very manner in which that person interprets and evaluates any individuating information about the applicant that is available. Thus, contrary to skeptics' claims, this position suggests that stereotypes can be *particularly* powerful in such hiring decisions.

Critics often claim that the IAT is best at predicting small effects or micro-behaviors in lab settings, and less useful for predicting behaviors that would be typically considered discrimination (Landy, 2008; Mitchell & Tetlock, 2006). This outdated belief has been addressed and qualified by recent evidence (Greenwald, 2008; Nosek, Greenwald, & Banaji, 2007). But the broader point is that, even if this claim were true, we would argue that discrimination in micro-behaviors (in lab

and/or non-lab contexts) should not be dismissed so casually. When aggregated over time, a career's worth of micro-actions can add up to much larger effects on a person's employment status (Reskin, 2002). Managers may want to be egalitarian and fair in their decision making, but often the conditions under which they are working involve time pressures and distractions. These are exactly the conditions that would impair an explicit commitment to nonprejudice and increase the likelihood that implicit biases and stereotypes would be influential (Chugh, 2004). Decades of psychological research suggests that interviewers can negatively influence the person being evaluated without being aware of doing so, i.e., they create a self-fulfilling prophecy (e.g., Word, Zanna, & Cooper, 1974). This process allows a manager to believe in his or her objectivity when evaluating candidates yet still make a decision that was influenced by implicit bias. Although we agree that the IAT can predict more than these spontaneous, automatic behaviors, we believe their power ought not to be dismissed so quickly.

Links between Implicit Bias and Behavior

Two other important common themes emerge from the set of articles in this special issue. One such theme is that the contributors to this issue routinely draw connections between the constructs that they study and real-world behaviors. Some have noted the trend in social psychology to focus more on self-report or reaction time measures when collecting data (Baumeister, Vohs, & Funder, 2007). As various papers in this special issue suggest, such data are valuable on their own, but they can be richer and more informative when used in conjunction with data on actual behavioral outcomes. A survey of the empirical work in this special issue reveals research on a range of behaviors from scheduling a job interview to prescribing a medical treatment to reacting to a criminal suspect. Such a strong effort to use a variety of measures to represent the complex forces at work in social situations where discrimination may occur is a significant development. Moreover, the results from behavioral studies complement those using more strictly laboratory-based research paradigms and bolster various claims about the relationship between stereotyping, prejudice, and discrimination. From a scientific standpoint, such a multimethod approach is encouraging because it increases our level of certainty about the connections between a mental process such as implicit bias and a behavioral outcome such as disparate treatment of a minority group member (e.g., Leslie et al., 2008). For psychology in general, such research responds to the call for more behavioral research to enrich the field (Baumeister et al., 2007).

Bridging the Divide between Individual and Aggregate Levels of Analysis

The other important common theme in this collection of articles is that they examine individual and aggregate-level accounts for the disparities in question.

Both levels of analysis are important to consider, as is the *relationship* between the aggregate and individual levels of analysis. As we discussed earlier, the debate over implicit bias often hinges on whether the aggregate-level research establishes prejudice as a causal force at the individual level of analysis (Tetlock & Mitchell, 2009). As the papers included in this special issue highlight, the connection between individual-level accounts and aggregate-level disparities in different social domains is complicated and often likely to be established indirectly and inferentially than by more direct empirical evidence. Kang (2012), for example, in thinking about these issues in a legal context, develops four "quadrants" that can be derived by crossing an axis of *specificity* (i.e., scientific and nonscientific facts can be specific or general), and an axis that involves a *temporal perspective* (i.e., an ex ante or future-oriented actions, or a ex post time orientation that emphasizes actions that already have taken place).

Kang's approach would suggest that there are two ways to account for discriminatory actions that allegedly have taken place: for example, to use the IAT as a "prejudice polygraph" to address these specific claims (an approach that even the developers of the IAT have not advocated) or to rely upon general scientific facts, presented in the form of a social framework analysis (Monahan, Walker, & Mitchell, 2008) to address specific discrimination claims (see Fiske & Borgida, 2008). When it comes to accounting for ex ante events, Kang similarly argues that science can be deployed in two different ways: to use, for example, the IAT as a selection tool to screen out potentially biased employees before discriminatory actions occur, or to use a foundational framework of consensus scientific facts to assess whether a proposed public policy should be adopted. More generally, Kang's analysis suggests that there are clearly different ways to think about the connection between aggregate-level disparities and the individual or psychological level of analysis. It is crucial to distinguish between the availability and applicability of specific versus general scientific facts, and whether such facts are being appropriately used to understand actions that have taken place or might take place in the future. In other words, Kang's "quadrants" approach is a more nuanced approach to thinking about the validity of individual-level accounts for aggregate-level disparities. Kang's analysis, along with related work by other legal scholars (e.g., Faigman, 2010; Mitchell, Monahan, & Walker, 2011), thus has tremendous potential to guide future empirical work.

Conclusion

While it is important to acknowledge the ongoing debate over the IAT's validity, we do not believe this controversy is sufficient reason to dismiss implicit bias as an account for real-world racial and gender disparities in various social contexts. The contributors to this special issue consider implicit bias and other accounts for disparities in legal, professional, or employment contexts, or in health

care contexts. The evidence that these papers present is often enough to rule out alternative hypotheses, but implicit bias remains as an account for disparities that, to varying degrees, fits with the available data. Moreover, implicit bias research is not synonymous with the IAT: evidence of bias emerges through a variety of research methods, including but not limited to the IAT. Taken together, the research included in this terrific special issue represents a strong body of evidence in support of the claim that implicit biases are contributing to an understanding of ongoing real-world disparities. As such, we believe that implicit bias research will continue to play a crucial role in understanding and hopefully reducing these aggregate-level disparities as they surface in employment, legal, and health care domains.

References

Amodio, D. M., & Devine, P. G. (2006). Stereotyping and evaluation in implicit race bias: Evidence for independent constructs and unique effects on behavior. *Journal of Personality and Social Psychology*, *91*, 652–661. doi: 10.1037/0022-3514.91.4.652.

Arkes, H., & Tetlock, P. E. (2004). Attributions of implicit prejudice, or "Would Jesse Jackson 'fail' the Implicit Association Test?" *Psychological Inquiry*, *15*, 257–278. doi:10.1207/s15327965pli1504_01.

Baumeister, R. F., Vohs, K. D., & Funder, D. C. (2007). Psychology as the science of self-reports and finger movements: What ever happened to actual behavior? *Perspectives on Psychological Science*, *2*, 396–405. doi:10.1111/j.1745-6916.2007.00051.x.

Bendick, M., & Nunes, A. P. (2012). Developing the research basis for controlling bias in hiring. *Journal of Social Issues*, *68*(2), 238–262. doi: 10.1111/j.1540-4560.2011.01747.x

Blanton, H., & Jaccard, J. (2006). Arbitrary metrics in psychology. *American Psychologist*, *61*, 27–41. doi:10.1037/0003-066X.61.1.27.

Borgida, E., Deason, G., Kim, A., & Fiske, S. T. (2008). Stereotyping research and employment discrimination: Time to see the forest for the trees. *Industrial and Organizational Psychology*, *1*, 405–408. doi: 10.1111/j.1754-9434.2008.00074.x.

Brondolo, E., Libretti, M., Riviera, L., & Walsemann, K. (2012). Racism and the development of social capital. *Journal of Social Issues*, *68*(2), 358–384. doi: 10.1111/j.1540-4560.2011.01752.x

Chugh, D. (2004). Societal and managerial implications of social cognition: Why milliseconds matter. *Social Justice Research*, *17*, 203–222. doi:10.1023/B:SORE.0000027410.26010.40.

Cikara, M., Rudman, L., & Fiske, S. T. (2012). Discrimination by a thousand cuts? Accounting for gender differences in top-ranked publication rates in social psychology. *Journal of Social Issues*, *68*(2), 263–285. doi: 10.1111/j.1540-4560.2011.01748.x

Correll, Park, J., 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*, 1314–1329. doi:10.1037/0022-3514.83.6.1314.

Correll, J., Park, B., Judd, C. M., Wittenbrink, B., Sadler, M. S., & Keesee, T. (2007). Across the thin blue line: Police officers and racial bias in the decision to shoot. *Journal of Personality and Social Psychology*, *92*, 1006–1023. doi: 10.1037/0022-3514.92.6.1006.

Eggly, S., Griggs, J. J., Orom, H., Penner, L. A., & Underwood, W. (2012). Life-threatening disparities: The treatment of Black and White cancer patients. *Journal of Social Issues*, *68*(2), 328–357. doi: 10.1111/j.1540-4560.2011.01751.x

Faigman, D. L. (2010). Evidentiary incommensurability: A preliminary exploration of the problem of reasoning from general scientific data to individualized legal decision-making. *Brooklyn Law School*, *75*, 1–22.

Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology*, *54*, 297–327. doi: 10.1146/annurev.psych.54.101601.14522.

Fiske, S. T., & Borgida, E. (2008). Providing expert knowledge for the legal system: Social science in an adversary context. *Annual Review of Law and Social Science*, *4*, 123–148.

Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression formation from category-based to individuating processes: Influences of information and motivation on attention and interpretation. In M. Zanna (Ed.), *Advances in experimental social psychology* (Vol. *23*, pp. 1–74). New York: Academic Press.

Glaser, J., & Knowles, E. D. (2008). Implicit motivation to control prejudice. *Journal of Experimental Social Psychology*, *44*, 164–172. doi: 10.1016/j.jesp.2007.01.002.

Green, A. R., Carney, D. R., Pallin, D. J., Ngo, L. H., Raymond, K. L., Iezzoni, L.I., & Banaji, M. R. (2007). Implicit bias among physicians and its prediction of thrombolysis decisions for Black and White patients. *Journal of General Internal Medicine*, *22*, 1231–1238. doi: 10.1007/s11606-007-0258-5.

Greenwald, A. G. (2008). Landy is correct: Stereotyping can be moderated by individuating the outgroup and by being accountable. *Industrial and Organizational Psychology*, *1*, 430–435. doi: 10.1111/j.1754-9434.2008.00082.x.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology*, *74*, 1464–1480. doi: 10.1037/0022-3514.74.6.1464.

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*, 17–41. doi: 10.1037/a0015575.

Jost, J. T., Rudman, L. A., Blair, I. V., Carney, D. R., Dasgupta, N., Glaser, J., & Hardin, C. D. (2009). The existence of implicit bias is beyond reasonable doubt: A refutation of ideological and methodological objections and executive summary of ten studies that no manager should ignore. *Research in Organizational Behavior*, *29*, 39–69. doi: 10.1016/j.riob.2009.10.001.

Kang, J. (2012). The missing quadrants of anti-discrimination: Going beyond the "prejudice polygraph". *Journal of Social Issues*, *68*(2), 314–327. doi: 10.1111/j.1540-4560.2011.01750.x

Kunda, Z., & Sinclair, L. (1999). Motivated reasoning with stereotypes: Activation, application and inhibition. *Psychological Inquiry*, *10*, 12–22. doi: 10.1207/s15327965pli1001_2.

Landy, F. (2008). Stereotypes, bias, and personnel decisions: Strange and stranger. *Industrial and Organizational Psychology*, *1*, 379–392. doi: 10.1111/j.1754-9434.2008.00071.x.

Leslie, L. M., King, E. B., Bradley, J. C., & Hebl, M. R. (2008). Triangulation across methodologies: All signs point to persistent stereotyping and discrimination in organizations. *Industrial and Organizational Psychology*, *1*, 399–404. doi: 10.1111/j.1754-9434.2008.00073.x.

Mendes, W. B., Gray, H. M., Mendoza-Denton, R., Major, B., & Epel, E. S. (2007). Why egalitarianism might be good for your health: Physiological thriving during stressful intergroup encounters. *Psychological Science*, *18*, 991–998. doi: 10.1111/j.1467-9280.2007.02014.x.

Mitchell, G., Monahan, J., & Walker, L. (2011). Beyond context: Social facts as case-specific evidence. *Emory Law Review*, *60*, 1109–1155.

Mitchell, G., & Tetlock, P. E. (2006). Antidiscrimination law and the perils of mindreading. *Ohio State Law Journal*, *67*, 1180–1192.

Monahan, J., Walker, L., & Mitchell, G. (2008). Contextual evidence of gender discrimination: The ascendance of "social frameworks." *Virginia Law Review*, *94*, 1715–1749.

Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2005). Understanding and using the Implicit Association Test: II. Method variables and construct validity. *Personality and Social Psychology Bulletin*, *31*, 166–180. doi: 10.1177/0146167204271418.

Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2007). The Implicit Association Test at age 7: A methodological and conceptual review. In J. A. Bargh (Ed.), *Social psychology and the unconscious: The automaticity of higher mental processes* (pp. 265–292). New York: Psychology Press.

Pager D. (2003). The mark of a criminal record. *American Journal of Sociology*, *108*, 937–975. doi: 10.1086/374403.

Pager, D., & Quillian, L. (2005). Walking the talk: What employers say versus what they do. *American Sociology Review*, *70*, 355–380. doi: 10.1177/000312240507000301.

Pager, D., & Western, B. (2012). Identifying discrimination at work: The use of field experiments. *Journal of Social Issues*, *68*(2), 221–237. doi: 10.1111/j.1540-4560.2011.01746.x

Penner, L. A., Albrecht, T. L., Orom, H. A., Coleman, D., & Underwood, W. III. (2009). The role of prejudice, stereotyping, and discrimination in Black-White health disparities in the United States. In J. F. Dovidio, M. Hewstone, P. Glick & V. M. Esses (Eds.), *Handbook of prejudice, stereotyping, and discrimination*. London: Sage.

Penner, L., Dovidio, J. F., West, T. V., Gaertner, S., Albrecht, T., Dailey, R. K., & Markova, T. (2010). Aversive racism and medical interactions with black patients: A field study. *Journal of Experimental Social Psychology*, *46*, 436–440. doi: 10.1016/j.jesp.2009.11.004.

Phelps, E. A., O'Connor, K. J., Cunningham, W. A., Gatenby, J. C., Funayama, E. S., Gore, J. C., & Banaji, M. R. (2000). Amygdala activation predicts performance on indirect measures of racial bias. *Journal of Cognitive Neuroscience*, *12*, 729–738. doi: 10.1162/089892900562552.

Reskin, B. (2002). Rethinking employment discrimination and its remedies. In M. F. Guillen, R. Collins, P. England, & M. Meyer (Eds.), *New directions in economic sociology* (pp. 218–244). New York: Sage.

Rooth, D. O. (2010). Automatic associations and discrimination in hiring: Real world evidence. *Labour Economics*, *17*, 523–534. doi:10.1016/j.labeco.2009.04.005.

Rudman, L. A. (2008). The validity of the Implicit Association Test is a scientific certainty. *Industrial and Organizational Psychology*, *1*, 426–429. doi:10.1111/j.1754-9434.2008.00081.x.

Rudman, L. A., & Lee, M. R. (2002). Implicit and explicit consequences of exposure to violent and misogynous rap music. *Group Processes and Intergroup Relations*, *5*, 133–150. doi:10.1177/1368430202005002541.

Sadler, M. S., Correll, J., Park, B., & Judd, C. M. (2012). The world is not black and white: Racial bias in the decision to shoot in a multiethnic context. *Journal of Social Issues*, *68*(2), 286–313. doi: 10.1111/j.1540-4560.2011.01749.x

Tetlock, P. E., & Arkes, H. (2004). The implicit-prejudice exchange: Islands of consensus in a sea of controversy. *Psychological Inquiry*, *15*, 311–321.

Lerner, J. S., & Tetlock, P. E. (1999). Accounting for the effects of accountability. *Psychological Bulletin*, *125*, 255–275. doi:10.1037/0033-2909.125.2.255.

Tetlock, P. E., Mitchell, G., & Murray, T. L. (2008). The challenge of debiasing personnel decisions: Avoiding both under-and overcorrection. *Industrial and Organizational Psychology: Perspectives on Science and Practice*, *1*(4), 439–443. doi:10.1111/j.1754-9434.2008.00084.x.

Tetlock, P. E., & Mitchell, G. (2009). Implicit prejudice and accountability systems: What must organizations do to prevent discrimination? *Research in Organizational Behavior*, *29*, 3–38. doi:10.1016/j.riob.2009.10.002.

Word, C. O., Zanna, M. P., & Cooper, J. (1974). The nonverbal mediation of self-fulfilling prophecies in interracial interaction. *Journal of Experimental Social Psychology*, *10*, 109–120. doi:10.1016/0022-1031(74)90059-6.

EMILY L. FISHER is now an Assistant Professor of psychology at Hobart and William Smith Colleges. She received her PhD in social psychology from the University of Minnesota, Twin Cities, in 2011, with an interdisciplinary specialty in political psychology. Her research focuses on intergroup relations, with a particular focus on the role of social capital in structuring intergroup attitudes.

EUGENE BORGIDA is a Professor of psychology and law, and Morse-Alumni Distinguished Professor of psychology at the University of Minnesota. He is a fellow of the APS and the APA, and Past President of SPSSI. He has served on the Board of Directors for the Association of Psychological Science (APS) and the Social Science Research Council (SSRC). Borgida's research interests include social cognition, attitudes and persuasion, psychology and law, and political psychology.