

the kinds of people who share these “individuating” properties. We have stereotypes about good-looking people, hipsters, short people, smiley people, good students, and so forth. Stereotypes can also enjoy Venn diagram relationships (or fail to in interesting ways); we may have a consensual stereotype for male and female, and a different unique stereotype for Asian male that does not merely overlap the male and Asian stereotypes. In other cases, we might lack a stereotype for the various salient group memberships, so an Asian male hipster might be seen as more Asian male, or more hipster, depending on the context. Or he might be seen as some overlapping combination of the two stereotypes. What this suggests is that Jussim’s Stereotype Rationality Hypothesis underestimates our need for stereotypes in social cognition. It is rational and reasonable for me to use reliable stereotypes in person perception, but I don’t need to jettison the stereotype when I gain additional relevant or even “definitive” information about the person, especially since that may just invoke another stereotype. Instead, our deliciously complex cognitive capacities use all the relevant information, allowing us to construct rich models of other people. The hippy girl might love steak, but only when the cow was grass fed.

That brings me to my final issue. Social cognition is a triangle between two individuals and a particular context. The context, which includes the goals of the perceiver, is an important variable when it comes to person perception. The same black youth may appear to be a good student sitting in a college classroom, and a worrying threat on the street—even to the same perceiver. The studies of racial stereotypes that ask teachers to make predictions of their students are limited to the triangle of teacher–student–classroom. The teacher’s stereotype of black youth may be accurate when it comes to her students, not only because she is familiar with them, but because she is likely motivated to see them in a particular way (part of the context). So it would be an overgeneralization to say that the teachers have accurate stereotypes of black youth. Rather, teachers may have accurate stereotypes of black youth *in the classroom*. When Republicans and Democrats use wildly inaccurate stereotypes about one another, their motivation to see one another in a certain way is also part of the context. If we care to examine the range of accuracy and inaccuracy in stereotypes (and I hope we do), we need to create a taxonomy of the varieties of stereotype types, one that reflects the order and breadth of the stereotypes themselves and the relationships between them, as well as the contexts in which they are created and used. We need a discipline of stereotype studies.

Are stereotypes accurate? A perspective from the cognitive science of concepts

doi:10.1017/S0140525X15002307, e3

Lin Bian^a and Andrei Cimpian^b

^aDepartment of Psychology, University of Illinois, Champaign, IL 61820. ^bDepartment of Psychology, New York University, New York, NY 10003.

linbian2@illinois.edu andrei.cimpian@nyu.edu
<http://www.psychology.illinois.edu/people/linbian2>
<http://cimpianlab.com>

Abstract: In his 2012 book, Jussim suggests that people’s beliefs about various groups (i.e., their stereotypes) are largely accurate. We unpack this claim using the distinction between *generic* and *statistical* beliefs—a distinction supported by extensive evidence in cognitive psychology, linguistics, and philosophy. Regardless of whether one understands stereotypes as generic or statistical beliefs about groups, skepticism remains about the rationality of social judgments.

Jussim is doing psychology a service by prompting careful thinking about a number of topics in his book *Social Perception and Social Reality* (Jussim 2012). We will focus our comments on his

arguments about stereotypes, which he defines as “beliefs about the attributes of social groups” (Précis target article, sect 8.2, para. 1). Going against the seeming consensus in social psychology, Jussim suggests that stereotypes are largely accurate. Here, we unpack this claim using a conceptual distinction from the cognitive psychology of concepts (used widely in linguistics and philosophy as well): namely, the distinction between *generic* and *statistical* beliefs about a category (e.g., Carlson & Pelletier 1995; Cimpian et al. 2010; Leslie et al. 2011; Prasada 2000). Attending to this distinction allows a more precise analysis of the claim that stereotypes are accurate—an analysis that ultimately undermines this claim. As we explain below, if one defines stereotypes as *generic* beliefs about groups, then the evidence Jussim presents (all of which pertains to people’s statistical estimates about various group attributes) is largely irrelevant to their accuracy. By virtue of their very structure, generic beliefs have only a weak relation to the statistical criteria that Jussim uses to define accuracy. On the other hand, if one defines stereotypes as *statistical* beliefs about groups, then one may no longer be speaking to the bulk of social judgments. The literature on concepts suggests it is generic—not statistical—beliefs that people use most readily when reasoning about categories and their members. Thus, regardless of how one unpacks Jussim’s claims on this topic, the accuracy of people’s judgments about groups is still in doubt. In what follows, we first outline the distinction between generic and statistical beliefs and then proceed to discuss its implications for Jussim’s arguments.

1. The distinction between generic and statistical beliefs about categories.

To begin, consider the statements below:

1. Fewer than 1% of mosquitoes carry the West Nile virus.
2. Mosquitoes carry the West Nile virus.
3. The majority of books are paperbacks.
4. Books are paperbacks.

Statements (1a) and (2a) are statistical: They express a belief about a certain number or proportion of the members of a category. Statements (1b) and (2b) are generic: They express a belief about the category as a whole rather than a specific number, quantity, or proportion. (An easy way to check the latter claim is to try out each of these statements as an answer to a “How many?” question. Only the statistical statements will sound appropriate.)

The fact that generic claims—and the beliefs they express—are not about numbers or quantities has a crucial consequence: It severs their truth conditions from the sort of statistical data that one could objectively measure in the world. In other words, whether people judge a generic belief about a category to be true does not straightforwardly depend on how many members of that category display the relevant attribute. This point is illustrated by the examples above. Both (1a) and (1b) are considered true: Although very few mosquitoes actually carry the West Nile virus, participants judge the generic claim (that MOSQUITOES,¹ as a category, carry the West Nile virus) to be true as well (e.g., Prasada et al. 2013). In contrast, even though (2a) is true—paperbacks are indeed very common—few believe that BOOKS, as a category, are paperbacks (i.e., [2b] is false). Notably, these are not isolated examples. The literature is replete with instances of generic claims that either are judged true despite unimpressive statistical evidence or judged false despite overwhelming numbers (e.g., Carlson & Pelletier 1995; Leslie 2007; 2008). In fact, the rules that govern which generic beliefs are deemed true and which are deemed false are so baroque and so divorced from the statistical facts that many linguists and philosophers have spent the better part of 40 years debating them (e.g., Carlson & Pelletier 1995; Lawler 1973; Leslie 2008).

Importantly, all of the foregoing applies to beliefs about social groups as well (e.g., Cimpian & Markman 2011; Cimpian et al. 2012; Gelman et al. 2004; Leslie 2008; in press; Prasada & Dillingham 2006; 2009; Rhodes et al. 2012). The distinction between statistical and generic beliefs is operative regardless whether

these beliefs concern mosquitoes, books, and other categories of non-human entities, or women, African Americans, Muslims, and other categories of humans. This means that generic beliefs about social groups, just like other generic beliefs, are typically removed from the underlying statistics. For example, more people hold the generic belief that MUSLIMS are terrorists than hold the generic belief that MUSLIMS are female (see Leslie, in press; Leslie et al. 2011). However, there are vastly more Muslims who are female than there are Muslims who are terrorists. Most of us would even be able to report these statistics, as Jussim's own data suggest; yet, awareness of the statistics has little bearing on endorsement of the respective generic beliefs. Again, this is not an isolated example. Compare, for instance, "ASIANS are really good at math" and "ASIANS are right-handed." Many more people would agree with the former generic claim than with the latter, while simultaneously being aware that the statistics go the opposite way.

In summary, people's beliefs about categories are of two types: generic and statistical. Although the accuracy of statistical beliefs depends solely on the data available in the world (e.g., how many Muslims are terrorists vs. women), the judged truth of generic beliefs does not. Rather, generic beliefs can be—and often are—largely discrepant with the reality on the ground.

2. Implications for the argument that stereotypes are accurate. We now go on to spell out the implications of this body of work for Jussim's argument. Regardless of which sort of belief (generic or statistical) he had in mind when claiming that stereotypes are largely accurate—and we will discuss each possibility in turn—the force of his argument is considerably weakened by attending to the evidence presented above.

2.1. Stereotypes as generic beliefs. Let's first assume that stereotypes are generic beliefs about groups. Based on our reading of the literature, this is how many social psychologists conceive of stereotypes, even though they understandably don't use the term *generic*. (Actually, at least one social psychologist we know of *did* use the term: Bob Abelson, whose research team published some fascinating work on "generic assertions" about social groups in the 1960s [e.g., Abelson & Kanouse 1966; Gilson & Abelson 1965].)

If stereotypes are generic beliefs, the evidence Jussim presents—all of which is about people's statistical estimates concerning group attributes—does not legitimize the claim that stereotypes are accurate. As explained above, generic beliefs depend only in a loose sense on the statistics available in the world. As a result, one cannot justifiably claim that generic beliefs are accurate, at least using the commonsense notion of accuracy that Jussim himself operates with. Is it accurate to believe—as most people do—that MOSQUITOES carry the West Nile virus but not that BOOKS are paperbacks? Is it accurate to believe that ASIANS are good at math but not right-handed? Sure, these beliefs may be "accurate" in the sense of being endorsed by many people, but agreement is a poor substitute for accuracy. The accuracy of a belief about a group is more legitimately assessed, as Jussim does, in terms of whether it matches the world statistically.

Based on these considerations, it seems reasonable to claim that, if stereotypes are generic beliefs, then they are often *inaccurate*—out of touch with the statistical reality. This is, of course, what many social psychologists have been claiming all along, and their claims appear justified under this definition of stereotypes. To further drive home this point, we briefly lay out four types of evidence suggesting considerable inaccuracy in people's (generic) stereotypes.

2.1.1. GENERIC BELIEFS ARE OFTEN ENDORSED ON THE BASIS OF SCANT STATISTICAL EVIDENCE

As already illustrated, many common generic beliefs are about attributes that are infrequent (e.g., MUSLIMS are terrorists). Notably, generic beliefs based on limited statistical evidence have also been observed in more controlled settings—for example, in laboratory studies where participants were given information about the prevalence of various traits in unfamiliar

categories and then tested for their endorsement of the corresponding generic beliefs (e.g., Brandone et al. 2015; Cimpian et al. 2010). Thus, several types of evidence (obtained with participants spanning the range from 4-year-olds to adults) suggest a disconnect between endorsement of generic beliefs and the underlying statistical facts.

2.1.2. GENERIC BELIEFS ARE RESISTANT TO COUNTEREVIDENCE

Related to the point about weak dependence on statistical evidence, once a generic belief is adopted, it is not easily falsified by exposure to evidence that contradicts it. The generic belief that MOSQUITOES carry the West Nile virus is not discarded as soon as a mosquito bite—or tens, or hundreds—fails to infect us. (The same goes for law-abiding Muslims and Asian people who aren't good at math.) Experimental work supports this conclusion as well. For example, 4-year-olds who first heard that PAGONS (an unfamiliar category) are friendly and were then shown a counterexample ended up generalizing this trait to novel pagons as frequently as children who did not see the counterexample (Chambers et al. 2008), which suggests that the counterexample had no effect on their endorsement of the generic belief.

2.1.3. GENERIC BELIEFS GIVE THE IMPRESSION OF STRONG STATISTICAL SUPPORT

Even though generic beliefs are often adopted on the basis of little statistical evidence, they nevertheless suggest—for example, when expressed in conversation—that the relevant attributes are almost always present (Brandone et al. 2015; Cimpian et al. 2010). For example, imagine a person who wasn't familiar with how the West Nile virus is transmitted. What would this person infer if they heard that MOSQUITOES carry it? Would they assume that fewer than 1% of mosquitoes in the affected areas are carriers, or perhaps that many more—even a majority—are? The evidence supports the latter possibility. In fact, most participants assume prevalence levels of greater than 90% when exposed to unfamiliar generic facts (Cimpian et al. 2010). There is thus a stark asymmetry at the core of generic beliefs: Although they are largely independent of the underlying statistics at the stage when they are initially formulated, they immediately take on the appearance of being rooted in strong statistical uniformities. For anyone who has little firsthand familiarity with the actual facts, this asymmetry can lead to largely mistaken impressions about the state of the world.

2.1.4. GENERIC BELIEFS ARE ACCOMPANIED BY MISLEADING EXPLANATORY INTUITIONS

Generic beliefs have strong explanatory overtones. Specifically, generic claims are consistently interpreted as conveying deep, inherent properties of the relevant categories (e.g., Cimpian & Cadena 2010; Cimpian & Erickson 2012; Cimpian & Markman, 2009, 2011; Gelman et al. 2010; Rhodes et al. 2012; see also Cimpian & Salomon 2014). Thus, when we are exposed to generic beliefs about, say, WOMEN being bad at math or AFRICAN AMERICANS being violent, we are seldom neutral as to the source of the attributes described. Rather, we implicitly adopt an explanatory perspective on these attributes, viewing them as core, non-accidental aspects of what the relevant groups are like deep down. To the extent that many group characteristics are not actually due to their members' biological makeup, this explanatory component of generic beliefs provides additional reasons to be suspicious of their match with the world.

In summary, if stereotypes are conceived as generic beliefs, then the evidence suggests they display a considerable amount of inaccuracy.

2.2. Stereotypes as statistical beliefs. What if we defined stereotypes as statistical beliefs instead? In this case, the evidence Jussim presents seems consistent with the idea that stereotypes are largely accurate. While that may be so, committing to a definition of stereotypes as statistical beliefs about groups may be problematic

for another reason: A recent (yet already widely replicated) finding in the literature on concepts suggests that people often have difficulty reasoning with – manipulating, basing inferences on, etc. – statistical knowledge about categories (e.g., Gelman et al. 2016; Hampton 2012; Hollander et al. 2002; Jönsson & Hampton 2006; Khemlani et al. 2012; Leslie et al. 2011; Leslie & Gelman 2012; Meyer et al. 2011). In many circumstances, people tend to fall back on using generic representations, consistent with theoretical arguments that such representations are an easy “default” when reasoning about categories (e.g., Cimpian & Erickson 2012; Gelman 2003; Leslie 2008). Thus, even if people are at some level attuned to the statistical distributions of various attributes across various groups, such statistical knowledge may ultimately be less influential than people’s generic beliefs about the same attributes.

We illustrate this point with data from Khemlani et al. (2012), who measured people’s expectations about the traits of unfamiliar category members – a ubiquitous type of social judgment (e.g., how strongly do you expect the next Asian person you’ll meet to be good at math?). Khemlani et al.’s goal was to compare the extent to which these expectations are rooted in participants’ statistical estimates (e.g., what percentage of Asian people do you think are good at math?) versus their generic beliefs (e.g., do you believe that ASIANS are good at math?).² The results highlighted the powerful influence of generic beliefs. Although participants’ statistical estimates did explain unique variance in their expectations about unfamiliar individuals, their endorsement of the relevant generic beliefs was considerably more predictive of these judgments (with an effect size that was 53% larger). Based on this and other similar evidence, we suggest that people’s awareness of the statistical distributions of various traits may be less important to their social judgments than their generic beliefs are. Further research testing this (admittedly bold) claim would be in order, however.

In summary, if stereotypes are conceived as statistical beliefs, they may not provide as much insight into people’s actual social judgments as one might expect.

3. Conclusion. Stereotypes are generic or statistical beliefs about the attributes of groups. If they are generic, they are likely not very accurate. If they are statistical, they may not be as influential as our (often inaccurate) generic beliefs about groups. Either way, one remains skeptical about the rationality of everyday social judgment.

ACKNOWLEDGMENTS

We are grateful to Joe Cimpian, David Yeager, Alice Eagly, and the members of the Cognitive Development Lab at the University of Illinois for helpful comments on previous drafts of this commentary. The writing of this commentary was supported in part by National Science Foundation grant BCS-1530669.

NOTES

1. For added clarity, we will occasionally use small caps to indicate when we intend to refer to categories.
2. Khemlani et al. (2012) didn’t include this specific item, but we use it here for consistency.

Trustworthiness perception at zero acquaintance: Consensus, accuracy, and prejudice

doi:10.1017/S0140525X15002319, e4

Jean-François Bonnefon,^a Astrid Hopfensitz,^{b, c} and Wim De Neys^d

^aToulouse School of Economics, Center for Research in Management, 31015 Toulouse, France; ^bToulouse School of Economics, GREMAQ, 31015 Toulouse, France; ^cInstitute for Advanced Study, 31015 Toulouse, France;

^dLaPsyDE (CNRS Unit 8240), Sorbonne – Paris Descartes University, 75005 Paris, France.

jfbonnefon@gmail.com hopfensitz@gmail.com wim.de-neys@parisdescartes.fr
<https://sites.google.com/site/jfbonnefon>
<http://hopfensitz.weebly.com>
<http://www.wdeney.org>

Abstract: Research on trustworthiness perception from faces has unfolded in a way that is strikingly reminiscent of Jussim’s narrative in his 2012 book. Jussim’s analysis warns us against overemphasizing evidence about prejudice over evidence about accuracy, when both are scant; and reminds us to hold all accounts to the same standards, whether they call on societal biases or true signals.

In the conclusion of his book (Jussim 2012), Jussim mentions recent lines of research on accuracy, and in particular the accuracy of judgments at zero acquaintance, formed from photos of strangers. This comment continues the discussion engaged in these final pages, extending Jussim’s argument to recent research on trustworthiness perception. Trustworthiness perceptions are especially interesting because they play a critical role in cooperation, which is itself at the very crossroad of current research in biology, economics, and psychology.

We highlight in this commentary that research on the perception of trustworthiness has unfolded in a way that is strikingly reminiscent of Jussim’s overall narrative. First, the bulk of this research has focused on consensus – that is, whether people agree about who looks trustworthy, regardless of whether they are correct or incorrect in this assessment. Second, research on trustworthiness perceptions has emphasized their potential for social injustice over their potential accuracy. Third, it has been speculated that the potential accuracy of trustworthiness perception may be due to self-fulfilling prophecies. Before we unpack each of these three points, we want to stress that our goal is not to argue that trustworthiness perceptions are fully accurate. We actually believe that the accuracy of trustworthiness perceptions is quite limited, and that they can have untoward social consequences. We also believe, though, that there is a kernel of accuracy in trustworthiness perceptions that is of broad and substantial theoretical interest. Accordingly, we wish for the field to give it full attention. As we will illustrate, this will require researchers to avoid several pitfalls vividly described by Jussim.

A large body of research has shown that people robustly agree on who looks trustworthy and who does not (Todorov et al. 2015b). However, studies that established this agreement were typically silent on its accuracy. For example, one paper showed that children as young as 3 or 4 rated the trustworthiness of unknown faces in a way that was consistent with adult ratings of the same faces (Cogsdill et al. 2014). However, because these faces were artificially constructed by a computer model, there was no objective measure of trustworthiness against which these judgments could be compared. This is also true of another striking study which showed that ultrafast trustworthiness ratings after 100 msec exposure to a face were highly correlated with judgments delivered after unrestricted time (Willis & Todorov 2006). Another paper showed that Americans and Japanese gave broadly consistent ratings of the trustworthiness of political candidates, based on their pictures (Rule et al. 2010). In this case the candidates were real persons, but the study did not attempt to correlate their actual behavior to the trust they inspired.

Obviously, the main difficulty in assessing the accuracy of trustworthiness perceptions is to obtain information about the individuals in the pictures, which can serve as a benchmark of trustworthiness. Recent research on trustworthiness perceptions started to offer at least two solutions to this challenge, one based on economic games conducted in the laboratory, and another based on naturalistic decision-making in the courtroom. The first line of research utilizes well-known behavioral economics