

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

protocols such as Public Good games, Prisoners' Dilemmas, and Trust games. The Trust game in particular is well suited to capture the accuracy of trustworthiness perceptions. One variant of this game involves two players, the Investor and the Trustee. The Investor is endowed with an initial sum of \$10, and can choose whether to keep that money or transfer it to the Trustee. The only information available to the Investor is a photo of the Trustee. If the Investor transfers the money, the Trustee receives the \$10 plus an additional \$20. The Trustee then decides whether to keep the whole \$30 or to split it equally with the Investor. The players cannot communicate, will not play a second round, and are completely informed about these rules and procedures. In sum, the Investor needs to decide whether to trust the other player to split the money, and the Trustee can decide whether to honor or to abuse this trust. Accordingly, the accuracy of trustworthiness perception can be measured by comparing the decisions of the Investor to the strategies of the Trustees: an Investor would demonstrate perfect accuracy by transferring money to all Trustees whose strategy is to split the money, and not transferring any money to Trustees whose strategy is to keep the money.

Several articles using this protocol showed that Investors did better than chance when deciding who to trust (e.g., Bonnefon et al. 2013; De Neys et al. 2013, 2015; Stirrat & Perrett 2010). However, it should be stressed that accuracy in economics games is quite limited, since Investors rarely make more than 55% correct decisions, where random decisions would be accurate 50% of the time. More importantly, this level of accuracy is only observed with cropped pictures that eliminate all but inner facial features, and disappears with full pictures showing hairstyle and clothing (Bonnefon et al. 2013). Furthermore, while Investors can show limited levels of accuracy when making their decisions about money transfers, they show no such accuracy when explicitly asked to rate the trustworthiness of the Trustees in the pictures (Bonnefon et al. 2013). In sum, this line of research has offered some evidence for the accuracy of trustworthiness detection, but also showed that it was limited in size and subject to strong contextual restrictions.

A second line of research has emphasized legal decision-making contexts, in which pictures depict individuals who stood accused, or were convicted of a crime. This line of research is promising because it can address both the accuracy of trustworthiness perceptions (do people who engage in criminal activities look untrustworthy?) and their potential for social injustice (do people who look untrustworthy receive harsher sentences?). The available evidence is scant on both fronts, though, because very few studies actually measured perceptions of trustworthiness, as opposed to, for example, perceptions of dangerousness. We know of two studies of accuracy, which showed that faces of criminals were judged as less trustworthy than faces of exemplary citizens, but once more to a small degree and subject to contextual restrictions (Porter et al. 2008; Rule et al. 2013). At the same time, we know of two articles documenting prejudiced legal decisions stemming from trustworthiness perceptions. First, Porter et al. (2010) showed that mock juries required less evidence to arrive at a guilty verdict when a defendant looked untrustworthy. Second and most recently, Wilson and Rule (2015) showed that convicted criminals who looked untrustworthy were more likely to have received a death sentence rather than a life sentence. In sum, there are few demonstrations yet that trustworthiness perceptions in legal decision-making contexts are either accurate or noxious. In such a situation, caution is required when characterizing our state of knowledge. We should be careful, for example, not to claim yet that this line of research has robustly established the unreliable and nefarious nature of trustworthiness perceptions (Bonnefon et al. 2015; Olivola et al. 2014, Todorov et al. 2015a). Jussim's analyses, though, warn us of the forceful pull to interpret scant evidence as definitive when it supports bias or prejudice, but weak when it supports accuracy.

In addition to warning us against such asymmetrical interpretations of the data, Jussim highlights problematic double standards

that can crop up in discussions of accuracy, and more specifically when explaining accuracy as the result of self-fulfilling prophecies. In a recent review of inferences from faces, Todorov et al. (2015b) suggested that accurate trustworthiness perceptions may not imply any biological link between morphology and behavior, but could instead reflect a self-fulfilling prophecy stemming from a societal bias. According to this account, people who have the sort of looks that societal biases associate with untrustworthiness would experience discrimination, and become less cooperative as a result, even though they were just as trustworthy to begin with. This is a promising line of thought, but one that Jussim reminds us to examine just as critically as its alternative. Consider for example these four propositions:

- Strong Consensus – Because of societal biases, people show strong agreement about who looks untrustworthy.
- Strong Prejudice – Because of societal biases, individuals who look untrustworthy suffer from discrimination.
- Self-fulfilling prophecy – Individuals who suffer from discrimination become less trustworthy as a result.
- No Accuracy – People who look untrustworthy are not actually untrustworthy.

Clearly, this set of propositions is inconsistent, so one proposition must be incorrect. If we adopt the societal bias narrative, that is, [Strong Consensus]+[Strong Prejudice]+[No Accuracy], we must conclude that [Self-fulfilling prophecy] is false. We are not arguing that this is the case: rather, we call attention to Jussim's warning to scrutinize claims about self-fulfilling prophecies to the same extent that we scrutinize claims about accuracy.

In sum, recent advances in the field of trustworthiness perception at zero acquaintance show striking similarities with the research reviewed by Jussim, even though the accuracy of trustworthiness perceptions is not nearly as high as (and much more fragile than) the accuracy of the judgments considered by Jussim. Nonetheless, Jussim's warnings apply well to this developing field of research: We must remain careful not to overemphasize evidence about prejudice over evidence about accuracy, when both are scant; and we must be careful to hold all accounts to the same standards of evidence, whether they call on societal biases or true morphological signals.

Perceptions versus interpretations, and domains for self-fulfilling prophecies

doi:10.1017/S0140525X15002320, e5

Jennifer Church

Department of Philosophy, Vassar College, Poughkeepsie, NY 12604.

church@vassar.edu <http://faculty.vassar.edu/church/>

Abstract: I suggest two ways in which Jussim's extensive discussion (in his 2012 book) could be enriched: first, by exploring the distinction between perceptual judgments and interpretive judgments; second, by considering the power of expectations to be self-fulfilling in the case of young children and the case of fragile egos.

Lee Jussim's book (Jussim 2012) argues – against current orthodoxy – that the judgments we make about the people we actually encounter are largely unaffected by faulty expectations we have regarding those people. Furthermore, he argues, cultural stereotypes are mostly correct; so even in cases where cultural expectations do affect the judgments we make about a particular individual, those expectations are more likely to increase rather than decrease our knowledge of that person.

Jussim's use of the term "perception" is very broad. Any judgment we make about someone that we observe is considered a "person perception." Such judgments can be more or less specific,