

Durham Research Online

Deposited in DRO:

23 October 2018

Version of attached file:

Accepted Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Puddifoot, Katherine (2017) 'Stereotyping : the multifactorial view.', *Philosophical topics.*, 45 (1). pp. 137-156.

Further information on publisher's website:

<https://doi.org/10.5840/philtopics20174518>

Publisher's copyright statement:

Additional information:

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

Stereotyping: The multifactorial view

Katherine Puddifoot

Abstract This paper proposes and defends the multifactorial view of stereotyping. According to this view, multiple factors determine whether or not any act of stereotyping increases the chance of an accurate judgement being made about an individual to whom the stereotype is applied. To support this conclusion, various features of acts of stereotyping that can determine the accuracy of stereotyping judgements are identified. The argument challenges two existing views that suggest that it is relatively easy for an act of stereotyping to increase the chance of an accurate judgement being made. In the process, it shows why stereotyping that associates Black people more strongly than White people with criminality in the United States cannot be defended, and actions to reduce the stereotyping criticised, on the basis that engaging in this form of stereotyping increases the chance of accurate judgements. As each of these important conclusions is supported by results from empirical psychology, the discussion exemplifies and vindicates the naturalistic approach to epistemology, according to which psychological findings provide an important contribution to understanding the epistemic standing of beliefs.

1. Introduction

Humans often engage in stereotyping. Stereotypes associating social groups with particular attributes influence judgements about individuals perceived to belong to the social groups. Discussion of stereotyping frequently focuses on how it is morally objectionable, but there is another important topic relating to stereotyping, that is, the epistemic standing of stereotyping, and, in particular, whether it leads to correct or incorrect judgements about individuals to whom stereotypes are applied.

Many authors discussing the latter topic have concentrated on how judgements that are formed in the process of stereotyping are likely to be erroneous, distorted and poorly supported by the evidence (e.g. Lippmann 1922; Allport 1954; Saul 2013). However, the dominant view in the psychology literature is that stereotypes can be accurate or inaccurate, which leaves open the possibility that stereotyping can increase the chance of an accurate judgement being made about individuals to whom a stereotype is applied.

This paper focuses on two positions that make claims about the conditions under which stereotyping can have this positive effect: What shall be called the *single factor view of stereotyping* and the *dual factor view of stereotyping*. The single factor view says that there is only one feature of an act of stereotyping that determines whether or not it increases the chance of an accurate judgement being made. The dual factor view says that there are two factors. I challenge both of these views, presenting a *multifactorial approach to stereotyping*, according to which numerous features of an act of stereotyping determine this. One consequence of the fact that multiple factors determine whether or not an act of stereotyping increases the chance of an accurate judgement being made is that stereotyping less frequently produces accurate judgements than the single factor and dual factor views suggest because there are more factors that need to be in place for an act of stereotyping to do so.

In the process of outlining the multiple features of an act of stereotyping that determine whether or not it increases the chance of an accurate judgement being made, I repeatedly discuss a specific type of stereotyping. This is what shall be called for the sake of the current discussion *race-crime stereotyping*:¹ the application of a stereotype that associates Black people more strongly than White people in the United States with certain crimes. This type of stereotyping is used as a case study to show various ways that an act of stereotyping can fail to increase the chance of an accurate judgement being made. The discussion will consequently highlight various ways that race-crime stereotyping fails to lead to accurate judgements, leading instead to erroneous judgements. The discussion will, in turn, provide a response to a criticism that could be levelled against actions to reduce race-crime stereotyping based on an argument presented by Tamar Szabo Gendler (2011), which suggests that race-crime stereotyping can bring epistemic benefits and actions to reduce the stereotyping can bring epistemic costs because not engaging in race-crime stereotyping involves base-rate neglect.

The current discussion both exemplifies and vindicates naturalized epistemology, as defended by, among others,² Alvin Goldman in *Epistemology and Cognition*. Naturalized epistemology is the view that epistemologists should utilize results from psychology in their epistemological work. For example, Goldman argues that epistemologists should use psychological findings to identify the cognitive processes that lead to the formation of beliefs and to establish whether those processes tend to lead to true or false beliefs.³ The current discussion uses results from psychology, alongside those from sociology, to identify the cognitive processes that lead to the formation of stereotyping judgements and to discover the conditions under which the social cognition involving stereotyping is likely to produce accurate and inaccurate judgements. In the process it achieves some important goals: (I) It challenges some false but intuitive views about stereotyping; (II) It provides a framework for analysing the epistemic standing of judgements that result from stereotyping; (III) It shows that race-crime stereotyping cannot be defended, and actions to reduce the stereotyping criticised, on the basis that the specific form of stereotyping tends to increase the chance of accurate judgements being made. As each of these goals is achieved due to the adoption of the naturalistic methodology, the paper vindicates the adoption of the methodology.

2. Stereotyping as truth-conducive

¹ This is not, of course, to suggest that there are no other forms of stereotyping connecting members of racial groups with criminal activity.

² See also, for example, Quine (1969); Stich (1990); Stein (1996); Nagel (2013); Kornblith (2014).

³ Goldman claims that results from psychology can be used to identify whether or not a cognitive process is reliable. The aim of the current paper is not to show that stereotyping is always, mostly or more often than not likely to produce true beliefs. It identifies a number of features of an act of stereotyping that can lead to false beliefs being produced, but without empirical work establishing whether or not these features are more often than not present it is not possible to draw a well-supported conclusion about whether stereotyping mostly or more often than not leads to error. However, although conclusions will not be drawn about the general reliability of stereotyping, it will be shown in this paper that under specific conditions, when certain features are present, acts of stereotyping are likely to produce inaccurate judgements. These observations are supported by results from empirical psychology revealing the nature of human cognitive processes. The current discussion therefore fits the general mould of naturalized epistemology as defended by Goldman and others.

In the recent philosophical (e.g. Beeghly 2015) and psychological (e.g. Jussim et al 2009, Jussim 2012) literature, it has been argued that stereotypes should be defined as accurate *or* inaccurate associations between individuals and attributes, where the associations are made due to the perceived social group membership of the individuals. Under this conception of stereotypes, they meet the following definition:

Stereotypes: mental states that associate members of certain social groups more strongly than others with particular attributes, in virtue of their perceived social group membership (see, e.g. Hilton & von Hippel 1994).

The association between Black people and criminality is a stereotype because it associates members of a particular social group *Black people* more strongly than members of other social groups with a particular attribute i.e. *criminality*. Meanwhile, stereotyping meets the following definition:

Stereotyping: making a judgment about an individual that is influenced by a mental state associating members of a group, to which that individual belongs, more strongly than members of other groups with particular attributes, in virtue of their perceived social group membership.

Under this general definition, it is within the realm of possibilities that acts of stereotyping can produce accurate judgements about individuals. Moreover, if a person engages in stereotyping when making a judgement about an individual they can consequently be more likely to make an accurate judgement than they would have been if they had not engaged in stereotyping. It is this possibility that is the focus of the current discussion: that the judgements constitutive of acts of stereotyping (also referred to in the current discussion as *stereotyping judgements*), e.g. the judgement *Tim is untrustworthy* based on the application of the stereotype that *politicians are untrustworthy* and the perception that Tim is a politician, are sometimes more likely to be accurate than alternative judgements that might have been made in the absence of the stereotyping.

There is an intuitive view, the *single factor view of stereotyping*, according to which judgements constitutive of acts of stereotyping are more likely to be accurate than alternative judgements that might have been made in the absence of the stereotyping as long as the stereotype that is applied reflects an aspect of the social reality. A stereotype reflects some aspect of social reality as long as there is a regularity found within society and the stereotype leads a person to respond in a way that reflects the regularity.⁴ The following example illustrates the intuitiveness of the single factor view. Jones harbours a stereotype associating men more strongly than women with an interest in football. This stereotype reflects the reality that in Jones' society significantly more men than women have an interest in football. She meets a heterosexual couple, applies the stereotype, and consequently judges that the man is more likely than the woman to have an interest in football. *Prima facie*, the

⁴ The idea that a stereotype reflects some aspect of social reality is left underspecified for the sake of the current discussion, reflecting the fuzziness of people's intuitions with regards to stereotyping. I would suggest that when people hold the view that any act of stereotyping involving a stereotype that reflects reality is likely to produce an accurate judgement, their understanding of what it is to reflect reality is vague, and they would accept that a stereotype reflects reality merely because they can find a regularity in society that roughly corresponds to the stereotype.

application of the stereotype will increase the chance of an accurate judgement being made about who will have an interest in football because it will lead to a judgement that is informed by information that is relevant and important relative to the judgement being made: information about the distribution of a particular trait (i.e. interest in football) across a society. The single factor view can be characterised in the following way to capture the intuitiveness of examples like these:

Single factor view: judgements constitutive of acts of stereotyping are more likely to be accurate than alternative judgements that might have been made in the absence of the stereotyping, because the stereotyping leads to a judgement that is informed by information that is important and relevant relative to the judgement, if and only if the stereotype reflects an aspect of social reality.

This definition of the single factor view can be used to provide a more detailed characterisation of how it is that the application of a stereotype seems to be able to increase the chance of an accurate judgement being made: the stereotype encodes important and relevant information and the information informs a judgement via an act of stereotyping.

On the single factor view, it will be relatively easy for stereotyping to increase the chance of an accurate judgement being made. Another view recently defended in the psychological literature also implies this. According to Lee Jussim and colleagues (2009; Jussim 2012), a stereotype is only required to have one feature for stereotyping judgements to be more likely to be accurate than the relevant alternative judgements that might have been made—the stereotype must be accurate. The stipulation that a stereotype must be accurate is stronger than the stipulation that a stereotype must reflect some aspect of social reality. A stereotype might reflect some aspect of reality, some regularity found within society, while also being properly called inaccurate. For example, a person might associate having an interest in football more strongly with men than women, this might reflect the reality that within a society more men than women are interested in football, but the stereotype might be inaccurate because it underestimates or overestimates the differing degree of interest in football across the genders. (See section 3 for further discussion.)

In addition to the stereotype being accurate, on Jussim and colleagues' view, when stereotyping judgements are more likely to be accurate than alternative judgements that might have been made in the absence of the stereotyping it will also be the case that the information available about the specific case to which the stereotype is applied (e.g. the particular individual member of a social group) is ambiguous. Jussim and colleagues therefore defend what shall here be called the dual factor view:

Dual factor view of stereotyping: judgements constitutive of acts of stereotyping are more likely to be accurate than alternative judgements that might have been made in the absence of the stereotyping if and only if (a) the stereotype that is applied is accurate and (b) only ambiguous information is available about the case to which the stereotype is applied.

The following example is presented in support of the dual factor view. Jane is a lifelong occupant of Alaska and Jan is a lifelong occupant of New York. Both say that it is cold where they are today. Jussim and colleagues (2009) claim that because the

information about the temperature is ambiguous, one should be influenced in one's judgement about which place is colder by background information about the average temperatures in Alaska and New York. Similarly, they claim, when one only has ambiguous information about a specific member of a social group, one should be informed by the distribution of traits across social groups when making judgements about the individual. In their view, accurate stereotypes encode this information, so they should be applied. If, on the other hand, only unambiguous information about the specific case is available, one should not engage in stereotyping.

The current discussion challenges both the single factor and the dual factor views. It does not deny that the features that they identify determine whether or not some stereotyping judgements are more likely to be accurate than alternative judgements that might have been made in the absence of the stereotyping. Instead, it shows that there are more features that determine whether or not an act of stereotyping produces accurate judgements than these accounts suggest. As a result, it cannot justifiably be assumed that stereotyping judgements are more likely to be accurate than alternative judgements that might have been made in the absence of stereotyping simply because the judgements are the result of the application of stereotypes that reflect some aspect of social reality or are accurate, and the case-specific information is ambiguous. Other factors also determine if the judgements are more likely to be accurate than potential alternatives. The position defended in this paper is therefore called the *multifactorial view* of stereotyping. The multifactorial view provides reason for thinking that stereotyping will produce accurate judgements less often than the single factor and dual factor views imply because there are more factors that need to be in place for accurate stereotyping judgements to be produced than these accounts suggest.

In particular, as well as the features outlined by both the single factor and dual factor views, according to the multifactorial view, the following factors can determine whether or not an act of stereotyping produces an accurate judgement: (i) whether or not the stereotype that is applied leads to judgements fitting with accurate statistical information; (ii) whether or not the stereotype that is applied is triggered when relevant; (iii) whether or not the information encoded in the stereotype is balanced appropriately against the case-specific information so that the application of the stereotype does not lead to the distortion of case-specific information. For the next three sections, then, the aim is to highlight the role of these factors in determining the accuracy of stereotyping judgements.

3. Stereotype accuracy

In this section, I emphasise that whether or not an act of stereotyping makes people respond in a way that is fitting with accurate statistical information—a feature that is not captured by the single factor view⁵—can determine whether or not an act of stereotyping produces an accurate judgement. It is possible to see how this feature can have this role by considering what happens when the feature is absent.

⁵ The dual factor view of stereotyping fares better than the single factor view. This is because it stipulates that a stereotype must be accurate rather than simply reflecting some aspect of social reality. The difference between the claim that a stereotype must reflect reality and the claim that it must be accurate is explored in section 2.

To understand how the fittingness of people's responses to accurate statistical information can determine the accuracy of stereotyping judgements it is valuable to consider the explanation, outlined in the previous section, of how stereotyping can increase the chance of an accurate judgement being made. The application of a stereotype can have this positive effect because stereotyping can produce judgements that are informed by important and relevant information. As stereotypes associate members of social groups with attributes in virtue of social group membership, the important and relevant information that they can supply is statistical information about the distribution of traits across a social group. For an act of stereotyping to increase the chance of an accurate judgement it should therefore lead a person to respond in a way that is fitting with accurate statistical information about the distribution of traits across a social group. If an act of stereotyping produces a judgement that is not fitting with accurate statistical information, for instance, leading a person to respond in a way that overestimates or underestimates the prevalence of an attribute in some social groups, then there is good reason to doubt that the stereotyping judgement that is produced will be accurate.

The literature on race-crime stereotyping shows that stereotypes can lead people to form judgements that are not fitting with accurate statistical information. In this literature, research has been undertaken to measure people's judgements about rates of criminal involvement and compare these judgements to crime statistics.⁶ A number of studies have found that people tend to overestimate the percentage of violent crimes committed by Black people (Chiricos et al 2004; Pickett et al 2012). People also tend to misjudge the types of crimes that Black people engage in: overestimating the extent to which Black individuals engage in motor vehicle theft, rape and criminal homicide (Gordon et al 1996). People have also been found to automatically associate Black people with drug crime—when asked to close their eyes and picture a drug user, 95% of survey respondents replied that they imagined a Black person (Watson Burston et al 1995, cited in Alexander 2011). However, the reality is that White and Black people engage in drug use at roughly the same rate in the US (Substance Abuse and Mental Health Services Administration 2004), and the use of some drugs such as heroin, cocaine, stimulants or methamphetamine is higher among the White population than the African American population (Office of Applied Studies, Substance Abuse and Mental Health Services Administration, 2007). Moreover, people tend to significantly overestimate the risk of crime in areas with a high Black population (Quillian and Pager 2010), suggesting that they associate the Black population with crime substantially more than they would if their stereotypes accurately reflected the statistical reality. Each of these studies shows how the phenomenon that is the focus of the current section can occur: a stereotype is applied that associates Black people with criminal activity, but the stereotype overestimates the rates of criminal involvement among the Black population, and, as a result, an act of stereotyping fails to increase the chance of an accurate judgement, instead increasing the chance of an inaccurate judgement.

⁶ There is a shortcoming of this methodology: it relies on arrest rates, which can fail to reflect the reality of criminal involvement (see, for example Alexander 2011 on the discrepancy between arrest rates and drug use across the races in the US). However, there is little reason to think that the people interviewed in these surveys had a more accurate source of information than crime rates on which they were basing their judgements. That their judgements are not reflecting crime rates suggests that they are not properly reflecting the best available evidence of this sort, which is, admittedly, imperfect.

The psychological and sociological literature provides explanations of why people who engage in stereotyping can fail to make judgements that are fitting with statistical realities. The first two features of human psychology that contribute to the explanation are that: (I) humans tend to remember extreme members of a group better than less extreme members (Rothbart et al 1978), and (II) what we remember determines how we stereotype a group (ibid). We are therefore predisposed to reflect people who possess extreme characteristics more heavily in our stereotypes than people who possess less extreme characteristics (ibid.). For instance, as criminals, all else being equal, possess more extreme characteristics than law-abiding people, we are disposed to reflect actions of the former more than those of the latter in our stereotypes of any group containing both (ibid.). This means that if we are exposed to descriptions of actions of members of a racial group, and some members of that group have engaged in crimes, the criminal activities are more heavily represented in the stereotypes formed about the group than the law-abiding activities. The likelihood that any member of the group will commit a crime will consequently be overestimated.

This phenomenon is combined with a number of social phenomena meaning that people are especially likely to develop and apply stereotypes that reflect a distorted picture of the characteristics of some social groups. The media serves as an important “cultural parent”, heavily influencing how members of social groups are viewed (Weisbuch et al 2009),⁷ and the general population of a society is sometimes exposed by the media to a disproportionately negative picture of members of a social group. For example, the US population is exposed to a disproportionately negative picture of Black involvement in crime via the media. With regard to crime in general, although the proportion of Black and White suspects depicted in the news media has been found to fairly accurately reflect the proportion of crimes committed by members of the two racial groups,⁸ there are far fewer positive Black role models (e.g. police officers, news reporters) depicted in the news media than there are positive White role models. This means that Black people are more heavily associated than White people with the negative rather than the positive aspects of crime.⁹ In addition to this, Black suspects are often depicted in a non-individualised way; they are more likely to be unnamed, and they are less likely to have an opportunity to speak or have a spokesperson speak for them in the media (Entman 1992). The media thereby fails to reflect the variation between individual members of the racial category, leading to homogenisation of Black people. There is also reason to think that news reports of some actions of Black people rely on stereotypes of criminality while reports of the same actions, when they are conducted by White people, do not.¹⁰

⁷ There is some reason for thinking that Black people’s stereotypes will be more accurate than White people’s, because Black people notice more variation between members of their own race (Judd & Park 1993) but also some studies suggesting that Black people’s stereotypes regarding crime are as inaccurate, if not more inaccurate, than White people’s (see, e.g. Sniderman and Piazza 1993).

⁸ Black people are depicted as suspects at a rate roughly fitting with the rate of arrests (Entman 1992, Chiricos & Escholz 2002).

⁹ For instance, a study of local news media in Orlando Florida found that only 1 in 20 White people portrayed in the news were crime suspects while 1 in 8 Black people portrayed in the news were suspects (Chiricos and Escholz 2002).

¹⁰ See, for example, how the action of taking some drinks from a shop in the aftermath of Hurricane Katrina was described in media depictions as ‘looting’ when undertaken by a Black person and ‘finding’ when undertaken by a White person: http://www.salon.com/2005/09/02/photo_controversy/. Another example is the differing media responses to these two separate cases of 7-year-old boys, one

The media influence combines with the psychological disposition to overestimate extreme examples in stereotypes, producing stereotypes that exaggerate the proportion of Black people who are involved in negative criminal activities. There is little reason to hope that this phenomenon is restricted to stereotypes about race and crime, so that the damage is limited and only race-crime stereotypes are consequently likely to lead to judgements that fail to accurately reflect the statistical reality. Wherever a disproportionately negative (or positive) media depiction is found, this can combine with the psychological disposition to overestimate extreme examples, leading to stereotypes that significantly exaggerate the proportion of a social group who have negative (or positive) characteristics.

The psychological and sociological literature therefore present very good reasons for thinking that judgements constitutive of acts of stereotyping can fail to be more accurate than alternative judgements that might have been made in the absence of the stereotyping because the stereotyping produces a judgement that is not fitting with accurate statistical information about the distribution of attributes across social groups. When an act of stereotyping leads to a judgement that is not fitting with accurate statistical information, it is inappropriate to describe it as leading to judgements that are informed by important and relevant information. As we have reason to think that some stereotyping judgements are likely to be accurate because they can reflect important and relevant information, evidence that a certain type of stereotyping does not lead to judgements that reflect important and relevant information constitutes evidence that it will not tend to produce accurate judgements. Therefore, acts of stereotyping that lead to judgements that are not fitting with accurate statistical information are highly unlikely to produce accurate judgements. We have now identified the first factor that can determine the accuracy of a stereotyping judgement that is not captured by both of the single and dual factor views of stereotyping (because it is not captured by the single factor view): whether or not the act of stereotyping leads to judgements fitting with accurate statistical information.

The advocate of the single factor view might object that their view does capture the feature of stereotyping discussed in this section. Recall that the single factor view stipulates that an act of stereotyping is more likely to be accurate than a judgement that would have been made in the absence of the stereotyping if and only if the stereotype that is applied reflects some aspect of social reality. It might be argued that if a stereotype reflects some aspect of social reality then it will lead to judgements that are fitting with accurate statistical information about the distribution of traits across social groups. The feature of stereotyping identified in this section would therefore be the same feature that is the focus of the single factor view. The mistaken nature of this response can be seen by returning to one of the stereotypes discussed in this section that associates Black people with criminal activities: the stereotype associating Black people more strongly than White people with drug crimes. The stereotype could accurately reflect an aspect of social reality, i.e. high *arrest rates* among the Black population for drug use, while also leading to judgements that fail to fit accurate statistical information about *actual* rates of drug use, which are similar across the

white and one Black, who each took a relative's car for a joyride:
<https://thesocietypages.org/socimages/2013/07/17/framing-childrens-deviance/>

Black and non-Black populations (Alexander 2011). In applying the stereotype associating Black people more strongly than White people with drug use, a person might therefore form a judgement reflecting an aspect of social reality. But they could nonetheless fail to accurately reflect the statistical reality of rates of drug use across different social groups. Their stereotyping judgement could then fail to fit accurate statistical information even though the stereotype, and the judgement influenced by the stereotype, reflects some aspect of social reality.

4. Stereotypes and relevance

Those who advocate the single factor of stereotyping might respond to section 3 that it is not surprising that the application of a stereotype does not increase the chance of an accurate judgement when the stereotype fails to make people respond in a way that is fitting with accurate statistical information. They might maintain, however, that whenever a stereotype leads people to respond in a way that is fitting with accurate statistical information the act of stereotyping will be likely to produce an accurate judgement. This would require only a minor modification to their view. Those who defend the dual factor view might highlight that their position is consistent with the content of section 4. They could point out that their position implies that an act of stereotyping will only be likely to produce an accurate judgement when the stereotype leads to judgements that are fitting with accurate statistical information, maintaining their original stance that a stereotyping judgement is likely to be accurate as long as it meets this criterion and case-specific information is ambiguous. This section shows that neither of these responses is satisfactory because there is another factor that can determine whether or not an act of stereotyping produces an accurate judgement: whether or not the stereotype is applied when relevant.¹¹

Research from social psychology provides strong support for the claim that individuals are not appropriately sensitive to contexts in which any statistical information that might be encoded in the stereotype is relevant to a judgement. Two triggers for stereotyping identified by psychologists are a wounded ego and the desire to justify the current social system (see, e.g. Fiske et al. 2007; Uhlmann et al. 2010), and there is good reason to doubt that either of these factors correlates with stereotypes being relevant to a judgement. This section shows how, as a consequence of these facts about human psychology, a stereotyping judgement can fail to be accurate, and an act of stereotyping can fail to increase the chance of an accurate judgement, because a stereotype is applied when irrelevant.

Consider first the fact that stereotypes are often triggered as a result of an individual suffering a wounded ego (Spencer et al 1998). Steven Fein and Steven Spencer (1997) found that people who were given poor feedback on an intelligence test were subsequently more likely to stereotype a gay man. Lisa Sinclair and Ziva Kunda (1999) found that if a person was criticised by a Black doctor they were subsequently more likely to engage in negative racial stereotyping. These and similar findings suggest that an individual's stereotypes are often activated as a result of a threat to his or her ego. Even if there are times when particular social stereotypes are relevant,

¹¹ Because the current discussion focuses on epistemic costs of stereotyping, the notion of relevance is one according to which a stereotype is relevant to a judgement if being influenced by the stereotype is likely to increase the chance of an individual achieving various distinctively epistemic goals, like knowledge or understanding, rather than other goals, e.g. wellbeing.

they are extremely unlikely to be precisely those times when an evaluator has a wounded ego. It is very difficult to imagine what would explain this correlation. On the contrary, it is not difficult to imagine a situation in which an evaluator has his ego wounded, a stereotype is activated, but the stereotype is irrelevant. Take the following example. A police officer approaches the car of a Black male, who has been pulled over for a minor traffic violation, e.g. one of his headlights is not working. The police officer asks the man to step out of the vehicle but he responds slowly to the command. The police officer is offended at what he takes to be a threat to his authority. This triggers a stereotype associating the innocent man with crime, the police officer evaluates the man as a criminal and treats him with hostility, and this leads to an escalation of tension and hostility between the two individuals. The stereotype associating Black people with crime is triggered although the Black man has not committed a crime, only a minor traffic violation. Examples like this illustrate that people's stereotypes, including but not exclusively their race-crime stereotypes, can be triggered when irrelevant.

A similar point can be made concerning the second factor thought to trigger the activation of stereotypes. Social psychologists believe that people's stereotypes are activated when they feel a threat to the current social system and are motivated to defend it. This hypothesis is taken to explain, for example, why stereotypes differ over time in a way that is consistent with defending the current system. For instance, it is taken to explain why while slavery was legal in the US, stereotypes of Black people were that they were happy, childlike and affectionate, but with the fight to abolish slavery this stereotype changed to represent Blacks as threatening (Uhlmann et al 2010; Alexander 2011). If stereotypes are activated by challenges to the status quo, and challenges to the status quo were correlated with the relevance of race-crime stereotypes, then there would be reason to think that the stereotypes are activated when relevant. However, just as in the case of the wounded ego, there is good reason to doubt that such a correlation exists. It is once again difficult to imagine why a correlation would exist but easy to imagine examples in which a stereotype is triggered, in this case due to a challenge to the current social system, but the stereotype is not relevant.

Let us reconsider race-crime stereotyping to see this point. Even if there are times when stereotypes associating Black people and crime are relevant, there is little reason to think that these are the same times when people feel a threat to the current social order and are motivated to defend it. On the contrary, it is not difficult to imagine a situation in which people are inclined to defend the current social order, negative stereotypes associating Black people with crime are consequently activated, and the activation of the stereotype is wholly inappropriate. Consider, for instance, a discussion between two individuals regarding changes to the US criminal justice system made under the "Fair Sentencing Act of 2010 (Public Law 111-220)". The bill reduced the difference in the sentences attached to the possession of crack and powder cocaine. The change to the law reduced discrimination against Black people in the legal system because Black people are statistically more likely than White people to be in possession of crack rather than powder cocaine, so were previously substantially more likely to receive the harsher sentence. Research suggests that when the issue was initially discussed, because the discussion would have involved scrutiny of the current social order, the judgements made by at least some of the discussants would have been heavily influenced by stereotypes that vindicate the status quo. The

stereotype strongly associating Black people with crime is likely to have been triggered, meaning that some participants in the discussion were inclined to defend the current high rates of incarceration among Black people. However, the stereotype is not relevant to this discussion, which should focus on whether specific sentences are appropriate for particular crimes. Here we therefore find reason to think that criminal stereotypes are triggered, due to the desire to defend the social order, where they are irrelevant.

What this discussion shows is that human psychology predisposes us to apply stereotypes when they are not relevant to the judgements we are making. If a stereotype is irrelevant then any information that is encoded in the stereotype will be irrelevant. The application of the stereotype will not lead to a judgement that is well-informed by important and relevant information. It will therefore not increase the chance of an accurate judgement being made about an individual case. On the contrary, the act of stereotyping will be likely to decrease the chance of an accurate judgement being made. We have therefore identified the second factor that can determine whether they produce accurate judgements that is not captured by both of the single and dual factor views (in this case it is captured by neither): (ii) whether the stereotype that is applied is triggered where relevant.¹²

5. Weighting of information

Let us now turn to a third factor that can determine whether or not an act of stereotyping produces an accurate stereotyping judgement that is not captured by the single or dual factor view: whether or not any background statistical information encoded in the stereotype is appropriately balanced against information about a specific case being judged. Take, for example, a judge in a criminal court making a judgement about a case. She might associate members of some racial groups more strongly than members of other racial groups with criminal activity. Let us suppose for the sake of discussion that the association encodes statistical information more accurately than psychological results suggest it would. She also applies the race-crime stereotype in a circumstance in which the stereotype is relevant. The application of a stereotype could nonetheless reduce the probability of an accurate judgement being made by reducing the extent to which the judgement on the case is consistent with important and relevant information. This is because it could prevent important and

¹² It might be objected that the dual factor view captures (ii). Recall that on the dual factor view, one factor that determines whether or not an act of stereotyping is likely to produce an accurate judgement is whether ambiguous or unambiguous case-specific information is available. One might think that stereotypes are only relevant if one lacks unambiguous case-specific information. Then the claim that the application of a stereotype increases the chance of an accurate judgement being made only if information encoded in the stereotype is relevant might seem to be equivalent to the claim that the application of a stereotype increases the chance of an accurate judgement being made only in cases where one lacks unambiguous information. However, it would be a mistake to conflate the relevance of the stereotype and the ambiguity of the case-specific information. A stereotype can fail to be relevant even in a situation in which the person applying the stereotype has only ambiguous case-specific information. Consider, for example, the situation in which a police officer asks the man to step out of the vehicle and applies the stereotype associating Black people with crime. In this case, the police officer might only have ambiguous evidence about whether the person in the car is a criminal, e.g. only the information that the person has committed a minor traffic violation. However, this does not mean that the stereotype associating Black people with criminality is relevant. The relevance of a stereotype is therefore clearly distinct from the ambiguity of the case-specific information.

relevant *case-specific* information from being properly accessed and processed. If the case-specific information is not properly accessed and processed it will not be balanced appropriately against any information encoded in the stereotype.

Psychological results point towards various ways that the application of a stereotype can prevent case-specific information from being properly accessed and processed. First of all, when a stereotype is activated, ambiguous behaviour can consequently be interpreted in a way that is fitting with the stereotype rather than as it should be viewed: i.e. as ambiguous (Duncan 1976).¹³ Second of all, only some relevant information—that which is consistent with the stereotype—is likely to be noticed while other relevant information, which is inconsistent with the stereotype, is not (e.g. Cohen 1981).¹⁴ Thirdly, when stereotypes are activated, people are often assumed to be more similar to members of their own social group (Bartsch and Judd 1993) and less similar to members of other social groups (e.g. Tajfel 1981) than they really are. Information about dissimilarities and similarities respectively can be missed. Finally, when stereotyping occurs, the people to whom a stereotype is applied can be given less opportunity than they should be to communicate their ideas.¹⁵ For example, Lisa Cooper and colleagues (2012) found that where physicians scored highly on measures of stereotyping of Black people as non-compliant, patients were less likely to be given the opportunity to speak and explain their symptoms. Information that people might supply, including information about their specific case that might otherwise be balanced against information encoded in the stereotype, can consequently be inaccessible as a result of stereotyping.

For example, in cases of race-crime stereotyping, these distortions of case-specific information can manifest in the following ways. If a judgement is made about the guilt or innocence of a Black suspect of a crime: (a) ambiguous behaviours of the suspect can be interpreted as consistent with criminality; (b) features of the suspect that fit the criminal stereotype can be noticed, attended to and remembered while others are ignored; (c) differences between the suspect and other previously encountered Black men can be missed; (d) similarities between the Black suspect and non-Black individuals who have positive features that suggest that they are not criminals can also be missed; (e) the Black suspect can be given less opportunity to explain his situation than if the stereotype was not applied. Where these phenomena occur, only distorted case-specific information will be available to be balanced against the stereotype.

On many occasions, case-specific information will be more diagnostic than any background statistical information that might be encoded in a stereotype. Under such circumstances, being influenced in one's judgement by a stereotype can reduce rather than enhance one's chance of making an accurate judgement that is well-informed by important and relevant information. This is because important and relevant case-specific information is not properly accessed and processed.

¹³ See Devine 1989 for a replication of this kind of result due to implicit stereotyping.

¹⁴ See Levinson 2007 for a replication of this kind of result due to implicit stereotyping.

¹⁵ Miranda Fricker's (2007) work on testimonial injustice and Kristie Dotson's (2011) work on testimonial smothering describe mechanisms via which people can be denied the opportunity to communicate their ideas, leading information that they might supply to not be accessible. It is beyond the scope of the current discussion to defend these views, but those committed to the views might be interested to note that the phenomena that they describe can play the distorting role outlined here.

It is worth considering a potential objection. Jussim and colleagues (2009; Jussim 2012) claim that people apply stereotypes exactly as they should do. They present evidence that where people possess unambiguous and complete case-specific information they do not apply stereotypes, only applying stereotypes where they possess ambiguous and incomplete evidence. Their argument might seem to establish that stereotypes and case-specific information are appropriately balanced. To see why this objection fails, consider the distinction between *information gathering* and *information possession*. Jussim and colleagues provide reason for thinking that people respond appropriately to information that they possess but do not consider how stereotypes prevent people from gathering unambiguous, complete case-specific information. The findings cited in this section show that there are many ways that the application of a stereotype can impede the information-gathering process.

6. Upshot of discussion so far

We are now in a position to see that there are more factors that can determine whether the application of a stereotype produces an accurate judgement about an individual case than the single and dual factor views imply. Because there are more factors that determine this than might be expected, there are also more ways that an act of stereotyping can fail to increase the chance of an accurate judgement being made relative to alternative judgements that might have been made in the absence of stereotyping. An act of stereotyping might not increase the chance of an accurate judgement being made even if it involves the application of a stereotype that reflects some aspect of social reality and/or is accurate and the person who engages in the act has only ambiguous case specific information. This could be either because the stereotype fails to make the person engaging in the stereotyping respond in a way that is fitting with accurate statistical information, because the stereotype is applied in an irrelevant context or because the application of the stereotype leads to the distortion of diagnostic case-specific information, preventing the information from being appropriately accessed and processed and balanced against the information encoded in the stereotype.

7. Race-crime stereotyping: the case-study

While the discussion in sections 3-5 has shown that it is more probable than might be expected that any particular act of stereotyping will fail to produce an accurate judgement, when it comes to one form of stereotyping—race-crime stereotyping—it is possible to draw an even stronger conclusion: that the stereotyping very often fails to produce an accurate judgement. Throughout sections 3 to 5 race-crime stereotyping was used to illustrate various factors that can determine whether or not an act of stereotyping produces an accurate judgement about an individual case, and in the process a picture has emerged according to which this specific form of stereotyping frequently does not do so. In section 3, evidence was introduced that suggests that race-crime stereotypes often fail to lead to judgements that are fitting with accurate statistical information about crime rates, and both psychological and sociological explanations of this phenomenon were given. In section 4, it was shown that race-crime stereotypes are often applied when they are not relevant, leading to judgements that are influenced by irrelevant information. In section 5, it was shown how the application of race-crime stereotypes can prevent important and relevant case-specific

information from being accessed and appropriately utilised. Overall, a picture has emerged according to which race-crime stereotyping frequently leads to errors in judgement rather than increasing the chance of an accurate judgement being made.

8. Race-crime stereotyping and base-rate neglect

In light of the observations in section 7, it is possible to respond to an argument presented by Tamar Szabo Gendler (2011, see also Egan 2011) that could be used to criticise actions aimed to reduce stereotyping.¹⁶ Gendler compares actions to reduce race-crime stereotyping to base-rate neglect. She claims that given the statistical reality that Black people are arrested at higher rates than White people for certain crimes, the rational agent ought to associate Black people more strongly than White people with those crimes. People who fail to make the association commit base-rate neglect, a phenomenon widely studied in the psychological literature in which participants ignore relevant background statistical information or base-rates when making a judgement. On this basis, Gendler claims that we face a dilemma with regards to stereotyping: we can either take action to prevent stereotyping, suffering the epistemic costs associated with base-rate neglect, “irrationality through base-rate neglect” (2011, 57), or we can suffer other epistemic costs that are associated with stereotyping that she has identified elsewhere in her work, “irrationality through encoding associations that you reflectively reject” (ibid.).

The phenomenon of base-rate neglect can be understood by considering the results of Tversky and Kahneman’s (1974) engineer-lawyer experiment. In the experiment, participants were told that 30 engineers and 70 lawyers had been interviewed and that short descriptions had been written about them based on the interviews. The participants were then given a description designed to be either more fitting of an engineer stereotype, more fitting of a lawyer stereotype, or to be neutral, that is, no more fitting of either stereotype. Participants were told that they had been randomly assigned one of these descriptions and were asked to rate the probability (on a scale of 0-100) that the description is of one of the engineers or one of the lawyers. When asked this question, participants should take the base-rate or background statistical information into account, i.e. the information about how many engineers and lawyers were interviewed. However, among participants who were given a neutral description, the median response was that the probability that the description was of one of the engineers or one of the lawyers was 50-50, suggesting that they did not take the base-rate information into account, instead judging each outcome to be equally likely because the description was neutral. Base-rate neglect can be viewed as a serious epistemic error because it involves ignoring important and relevant background statistical information when making a judgement.

By comparing actions to prevent race-crime stereotyping to base-rate neglect, Gendler suggests that they cause a similar epistemic error. We are now in a position to understand why this conclusion should be resisted. The empirical results discussed in this paper suggest that when people engage in race-crime stereotyping they depart significantly from the ideal of base-rate use. In ideal base-rate use, accurate background statistical information is available for use in a judgement, e.g. information about what percentage of interviewees were engineers and lawyers or

¹⁶ For another, compatible response to Gendler (2011), see Madva (2016).

information about what percentage of instances of crime were committed by people of a particular racial group. Then the person utilising the base-rate information is sensitive to whether the information is relevant to the particular context, only applying the information when relevant. Then, once a situation and some base-rate information have been paired up, the base-rate information (e.g. about numbers of engineers and lawyers interviewed) is balanced appropriately against other, relevant case-specific information (e.g. details of the description given). Neither the case-specific information nor the base-rate information is given too much or too little weight. Race-crime stereotyping often differs significantly from this ideal. It does not involve stereotypes that encode accurate base-rate information making it available to influence a judgement. It does not involve stereotypes consistently only being applied when they are relevant. Nor does it involve base-rates being appropriately balanced against case-specific information. Actions to prevent race-crime stereotyping consequently often do not prevent people from properly utilising base-rate information. It is therefore misleading to compare these actions to base-rate neglect and suggest that they similarly involve an epistemic error. It would be less misleading to emphasise how ordinary acts of race-crime stereotyping frequently differ significantly from appropriate use of base-rates and actions to prevent stereotyping stop people from acting in ways that are far from ideal.

9. Conclusions

There is a single factor view of stereotyping according to which judgements constitutive of acts of stereotyping are more likely to be accurate than alternative judgements that might have been made in the absence of the stereotyping as long as the stereotype reflects an aspect of the social reality. Furthermore, there is a dual factor view of stereotyping according to which judgements constitutive of acts of stereotyping are more likely to be accurate than alternative judgements that might have been made in the absence of the stereotyping if and only if the stereotype that is applied is accurate and only ambiguous information is available about the case to which the stereotype is applied. Neither of these views of stereotyping is correct. Instead, a multifactorial view of stereotyping is correct. There are more features of any particular act of stereotyping that determines if it increases the chance of an accurate judgement being made than the single and dual factor views imply. The features that determine this include: (i) whether or not the stereotype that is applied leads to judgements fitting with accurate statistical information; (ii) whether or not the stereotype is applied in a scenario in which case specific information is ambiguous; (iii) whether the stereotype that is applied is triggered when relevant; (iv) whether the information encoded in the stereotype is balanced appropriately against case-specific information or the application of the stereotype leads to the distortion of case-specific information. Any analysis of whether a particular act of stereotyping is likely to produce an accurate judgement should involve an examination of whether each of (i)-(iv) are present. The first two features are captured by the dual factor view, but the latter two are omitted from both the single factor and dual factor views.

With respect to race-crime stereotyping, there is good reason to believe at least (i), (iii) and (iv) are frequently not present. It is often the case that race-crime stereotypes fail to encode accurate statistical information, they are applied when irrelevant and lead to the distortion of case-specific information, so that diagnostic information about the specific case is not appropriately accessed and utilised. Because of this, it is

possible to respond to a criticism of actions to prevent race-crime stereotyping that finds support in recent work by Tamar Szabo Gendler (2011): that these actions lead to base-rate neglect. What has become clear in the discussion in this paper is that race-crime stereotyping differs significantly from ideal use of base-rates. It is therefore misleading to characterise acts to prevent stereotyping as preventing ideal base-rate use. Actions to prevent race-crime stereotyping are more appropriately characterised as preventing people from acting in a way that is far from ideal.

The methodology adopted in this paper has been to draw on results from empirical psychology to identify various factors that can determine whether acts of stereotyping increase the chance of accurate judgements being made by ordinary human reasoners. In the process, an intuitive picture of stereotyping, according to which stereotyping increases the chance of an accurate judgement being made as long as the stereotype that is applied reflects some aspect of social reality, has been challenged. Moreover, a framework has been developed for assessing the value of any particular act of stereotyping, with important practical applications, as illustrated by its application to race-crime stereotyping. The methodology adopted in this paper has thereby yielded valuable results. The current discussion therefore provides a vindication of naturalized epistemology, with its commitment to the idea that results from empirical psychology provide important insights about the epistemic standing of our beliefs.

Reference

- Alexander, M. (2011). "The New Jim Crow" *Ohio State Journal of Criminal Law*, 9-27.
- Allport, G. W. (1954). *The Nature of Prejudice* (New York: Perseus Books).
- Bartsch RA, Judd CM. (1993). "Majority-minority status and perceived ingroup variability revisited", *European Journal of Social Psychology*, 23: 471–83.
- Beeghly, E. (2015). "What is a stereotype? What is stereotyping?", *Hypatia*, 30, 4, 675-691.
- Chiricos, T., & Eschholz, S. (2002). "The racial and ethnic typification of crime and the criminal typification of race and ethnicity in local television news", *Journal of Research in Crime and Delinquency*, 39, 400-420.
- Chiricos, T., Welch K., & Gertz M. (2004). "The racial typification of crime and support for punitive measures", *Criminology*, 42, 358-390.
- Cohen, C. E. (1981). "Personal categories and social perception: Testing some boundaries of the processing effects of prior knowledge", *Journal of Personality and Social Psychology*, 40, 441-452.
- Cooper, L. A., Roter, D. L., Carson, K. A., Beach, M. C., Sabin, J. A., Greenwald, A. G., & Inui, T. S. (2012). "The associations of clinicians' implicit attitudes about race with medical visit communication and patient ratings of interpersonal care", *American journal of public health*, 102, 5, 979-987.

- Devine, P. (1989). "Stereotypes and prejudice: Their automatic and controlled components", *Journal of personality and social psychology*, 56, 1, 5.
- Dotson, K. (2011). "Tracking epistemic violence, tracking patterns of silencing", *Hypatia*, 26, 2, 236–257.
- Duncan, B. L. (1976). "Differential social perception and attribution of intergroup violence: Testing the lower limits of stereotyping of Blacks", *Journal of Personality and Social Psychology*, 34, 590-598.
- Egan, A. (2011). "Comments on Gendler, 'On the Epistemic Costs of Implicit Bias'", *Philosophical Studies*, 156, 1, 65-79.
- Entman, R. M. (1992). "Blacks in the news: Television, modern racism and cultural change". *Journalism Quarterly*, 69, 341-361.
- Fein, S., & Spencer, S. J. (1997). "Prejudice as self-image maintenance: Affirming the self through derogating others", *Journal of Personality and Social Psychology*, 73, 31–44.
- Fiske, S., Cuddy, A.J.C. & Glick, P. (2007). "Universal Dimensions of Social Cognition: warmth and competence", *Trends in Cognitive Sciences*, 11, 2, 77-83.
- Fricker, M. (2007) *Epistemic Injustice: Power and the Ethics of Knowing* (New York: Oxford University Press).
- Gendler, T. S. (2011). "On the Epistemic Costs of Implicit Bias", *Philosophical Studies*, 156, 1, 33-63.
- Goldman, A. (1986). *Epistemology and cognition*. (Harvard University Press).
- Gordon R. A., Michels, J.L. & Nelson C.L. (1996). "Majority Group Perceptions of Criminal Behavior: The accuracy of race-related crime stereotypes", *Journal of Applied Social Psychology*, 26, 2, 148-159.
- Hilton, J. L. and von Hippel, W. (1996). "Stereotypes", *Annual Review of Psychology*, 47, 237–71.
- Judd, C.M. & Park, B. (1993). "Definition and Assessment of Accuracy in Social Stereotypes", *Psychological Review*, 109, 110–11.
- Jussim, L. (2012). *Social Perception and Social Reality* (Oxford: Oxford University Press).
- Jussim, L., Cain, T. R., Crawford, J. T., Harber, K., & Cohen, F. (2009). "The unbearable accuracy of stereotypes", in T.D. Nelson (ed.) *Handbook of prejudice, stereotyping, and discrimination*, (New York: Psychology Press), 199-227.
- Kornblith, H. (2014). *A Naturalistic Epistemology: Selected Papers* (Oxford: OUP).

Levinson, J. D. (2007). "Forgotten racial equality: Implicit bias, decisionmaking, and misremembering", *Duke Law Journal*, 345-424.

Lippman, W. (1922). "Stereotypes". *Public opinion* (New York, NY, US: MacMillan), 79-94.

Madva, A. (2016). "Virtue, Social Knowledge and Implicit Bias", M. Brownstein & J. Saul (eds.) *Implicit Bias and Philosophy, volume 1: Metaphysics and Epistemology* (Oxford: OUP).

Nagel, J. (2013). "Knowledge as a Mental State", *Oxford Studies in Epistemology* 4, 273-308.

Office of Applied Studies, Substance Abuse and Mental Health Services Administration (2007). "National Survey on Drug Use and Health: Demographic and Geographic Variations in Injection Drug Use", *US Department of Health and Human Services; Washington, DC*.

Pickett, J.T., Chiricos, T., Golden, K.M., & Gertz, M. (2012). "Reconsidering the relationship between perceived neighbourhood racial composition and Whites' perceptions of victimization risk: Do racial stereotypes matter?", *Criminology*, 50, 145-186

Quillian, L. & Pager, D. (2010). "Estimating risk: Stereotype amplification and the perceived risk of criminal victimization", *Social Psychology Quarterly* 73, 79-104.

Quine, W.V.O. (1969). "Epistemology Naturalized", In *Ontological Relativity and Other Essays* (New York: Columbia University Press).

Rothbart, M., Fulero, S., Jensen, C., Howard, J & Birrell, P. (1978). "From Individual to Group Impressions: Availability Heuristics in Stereotype Formation", *Journal of Experimental Social Psychology*, 14, 3, 237-255.

Saul, J. (2013). "Scepticism and Implicit Bias", *Disputatio*, 5, 37, 243-263.

Sinclair, L., & Kunda, Z. (1999). "Reactions to a Black professional: Motivated inhibition and activation of conflicting stereotypes", *Journal of Personality and Social Psychology*, 77, 885-904.

Sniderman, P., & Piazza, T. (1993). *The scar of race* (Cambridge, MA: Harvard University Press).

Spencer, S. J., Fein, S., Wolfe, C. T., Fong, C., & Dunn, M. A. (1998). "Automatic activation of stereotypes: The role of self-image threat", *Personality and Social Psychology Bulletin*, 24, 1139-1152.

Stein, E. (1996) *Without Good Reason: The rationality debate in philosophy and cognitive science* (New York: Oxford University Press).

Stich, S. (1990) *The Fragmentation of Reason*, (Cambridge, MA: MIT Press).

Substance Abuse and Mental Health Services Administration (2004) “Results from the 2003 National Survey on Drug Use and Health (NSDUH): national findings”, Available at: <http://store.samhsa.gov/product/Results-from-the-2003-National-Survey-on-Drug-Use-and-Health-NSDUH-National-Findings/BKOAS05>.

Tajfel, H. (1981). *Human Groups and Social Categories: Studies in Social Psychology* (Cambridge: Cambridge University Press).

Tversky, A. & Kahnemann, D. (1974). “Judgment under Uncertainty: Heuristics and Biases”, *Science, New Series*, 185, 4157, 1124-1131.

Uhlmann, E.L. Brescoll, V.L. & Machery, E. (2010). “The Motives Underlying Stereotype-Based Discrimination Against Members of Stigmatized Groups”, *Social Justice Research*, 23:1-16.

Watson Burston, B., Jones, D. & Robinson-Saunders, P. (1995). “Drug Use and African Americans: Myth Versus Reality”, *Journal of Alcohol and Drug Education*, 40, 2, 19-39.

Weisbuch, M., Pauker, K. & Ambady, N. (2009). “The subtle transmission of race bias via televised non-verbal behaviour”, *Science*, 326, 1711-1714.