ARTICLE

Implicit stereotypes and the predictive brain: cognition and culture in "biased" person perception

Perry Hinton[1]

ABSTRACT   Over the last 30 years there has been growing research into the concept of implicit stereotypes. Particularly using the Implicit Associations Test, it has been demonstrated that experimental participants show a response bias in support of a stereotypical association, such as "young" and "good" (and "old" and "bad") indicating evidence of an implicit age stereotype. This has been found even for people who consciously reject the use of such stereotypes, and seek to be fair in their judgement of other people. This finding has been interpreted as a "cognitive bias", implying an implicit prejudice within the individual. This article challenges that view: it is argued that implicit stereotypical associations (like any other implicit associations) have developed through the ordinary working of "the predictive brain". The predictive brain is assumed to operate through Bayesian principles, developing associations through experience of their prevalence in the social world of the perceiver. If the predictive brain were to sample randomly or comprehensively then stereotypical associations would not be picked up if they did not represent the state of the world. However, people are born into culture, and communicate within social networks. Thus, the implicit stereotypical associations picked up by an individual do not reflect a cognitive bias but the associations prevalent within their culture—evidence of "culture in mind". Therefore to understand implicit stereotypes, research should examine more closely the way associations are communicated within social networks rather than focusing exclusively on an implied cognitive bias of the individual.

[1] University of Warwick, Centre for Applied Linguistics, Coventry, UK Correspondence: (e-mail: p.r.hinton@warwick.ac.uk)

Traditionally a stereotype has been defined as overgeneralized attributes associated with the members of a social group (such as the reserved English or the geeky engineer), with the implication that it applies to all group members (Hinton, 2000). A large body of research, particularly in the United States of America (USA), has focused on the (negative) stereotypes of women and African Americans, which are linked to prejudice and discrimination in society (Nelson, 2009, Steele, 2010). Psychological researchers have sought to identify why certain people employed stereotypes and, in much of the twentieth century, they were viewed as due to a mental fallacy or misconception of a social group, an individual's "biased" cognition, resulting from proposed factors such as "simplicity" of thought (Koenig and King, 1964) and arising from upbringing and social motivation (particularly "authoritarianism", Adorno et al., 1950). A considerable amount of effort has been made subsequently to persuade people to avoid stereotype use, by highlighting its inaccuracy and unfairness (for example, Brown, 1965). However, since the 1960s, cognitive researchers, such as Tajfel (1969), have argued that stereotyping is a general feature of human social categorization. Despite this, it has been argued that individuals can consciously seek to avoid using negative stereotypes and maintain a non-prejudiced view of others (Devine, 1989; Schneider, 2004). Indeed, Fiske and Taylor (2013) claim that now only ten percent of the population (in Western democracies) employ overt stereotypes. Unfortunately, recent work, specifically using techniques such as the Implicit Associations Test (Greenwald et al., 1998), has shown that stereotypical associations can implicitly influence social judgement, even for people who consciously seek to avoid their use (Lai et al., 2016). These implicit stereotypes have provoked questions of both the control of, and an individual's responsibility for, the implicit effects of stereotypes that they consciously reject (Krieger and Fiske, 2006). This article explores the nature of implicit stereotypes by examining what is meant by "bias" in the psychological literature on stereotyping, and proposes an explanation of how culture influences implicit cognition through the concept of the "predictive brain" (Clark, 2013). The present work argues that, rather than viewing implicit stereotypes as a problem of the cognitive bias of the individual (for example, Fiske and Taylor, 2013), they should be viewed as "culture in mind" influencing the cognition of cultural group members. It is also proposed that combining the research on implicit cognition with an understanding of the complex dynamics of culture and communication, will lead to greater insight into the nature of implicit stereotypes.

Implicit stereotypes

The view of a stereotype as a fixed set of attributes associated with a social group comes from the seminal experimental psychology research by Katz and Braly (1933). One hundred students of Princeton University were asked to select the attributes that they associated with ten specific nationalities, ethnic and religious groups from a list of 84 characteristics. The researchers then compiled the attributes most commonly associated with each group. Katz and Braly (1933: 289) referred to these associations as "a group fallacy attitude", implying a mistaken belief (or attitude) on behalf of the participants. The study was repeated in Princeton by Gilbert (1951) and Karlins et al. (1969), and similar attributes tended to emerge as the most frequent for the groups. The endurance of these associations, such as the English as tradition-loving and conservative, over 35 years has often been narrowly interpreted as evidence for the fixed nature of stereotypes. Yet, a closer look at the data shows counter-evidence. Rarely was an attribute selected by more than half the participants: for the English only "sportsmanlike" in 1933, and "conservative" in 1969

reached this figure. Also both the percentages and the chosen attributes changed over time. By 1969, "sportsmanlike" for the English had dropped to 22%. A number of attributes in the initial top five for some of the groups dropped to below 10% by 1969. Also the stereotypes generally tended to become more positive over time. However, what the studies did establish was a methodological approach to stereotypes as the experimental investigation of "character" attributes associated with social groups in the mind of an individual.

The notion of implicit stereotypes is built on two key theoretical concepts: associative networks in semantic (knowledge) memory and automatic activation. Concepts in semantic memory are assumed to be linked together in terms of an associative network, with associated concepts having stronger links, or are closer together, than unrelated concepts (Collins and Loftus, 1975). Thus "doctor" has a stronger link to "nurse" (or viewed as closer in the network) than to unrelated concepts, such as "ship" or "tree". Related concepts cluster together, such as hospital, doctor, nurse, patient, ward, orderly, operating theatre, and so forth, in a local network (Payne and Cameron, 2013) that is sometimes referred to as a schema (Ghosh and Gilboa, 2014; see Hinton, 2016). Activation of one concept (such as reading the word "doctor") spreads to associated concepts in the network (such as "nurse") making them more easily accessible during the activation period. Evidence for the associative network model comes from response times in a number of research paradigms, such as word recognition, lexical decision and priming tasks: for example, Neely (1977) showed that the word "nurse" was recognized quicker in a reaction time task following the word "doctor" than when preceded by a neutral prime (such as a row of X's) or an unrelated prime word (such as "table"). Considerable amount of research has been undertaken on the nature of semantic association, which reflects subjective experience as well as linguistic similarity, although people appear to organize their semantic knowledge in similar ways to others. Weakly associated concepts may be activated by spreading activation based on thematic association, and the complexity of the structure of associations develops over time and experience (De Deyne et al., 2016).

The spreading activation of one concept to another was viewed as occurring unconsciously or automatically. In the mid-1970s a distinction was made between two forms of mental processing: conscious (or controlled) processing and automatic processing (Shiffrin and Schneider, 1977). Conscious processing involves attentional resources and can be employed flexibly and deal with novelty. However, it requires motivation and takes time to operate, which can lead to relatively slow serial processing of information. Automatic processing operates outside of attention, occurs rapidly and involves parallel processing. However, it tends to be inflexible and (to a high degree) uncontrollable. Kahneman (2011) refers to these as System 2 and System 1, respectively. Shiffrin and Schneider (1977) found that detecting a letter among numbers could be undertaken rapidly and effortlessly, implying the automatic detection of the categorical differences of letters and numbers. Detecting items from a group of target letters among a second group of background letters took time and concentration, requiring (conscious) attentional processing. However, novel associations (of certain letters as targets and other letters as background) could be learnt by extensive practice as long as the associations were consistent (targets were never used as background letters). After many thousands of trials, detection times reduced significantly, with the participants reporting the targets "popping out" from the background letters, implying that practice had led to automatic activation of the target letters (based on the new target-background letter categories). Thus, consistency of experience (practice) can lead

to new automatically activated learnt associations. However, when Shiffrin and Schneider (1977) switched the targets and background letters after thousands of consistent trials, performance dropped to well below the initial levels—detection times were extremely slow requiring conscious attention as participants struggled with the automatic activation of the old-but-now-incorrect targets. Slowly, and with additional practice of thousands of trials, performance gradually improved with the new configuration of target and background letters. Thus, highly practiced semantic associations—consistent in a person's experience—can become automatically activated on category detection —but once learnt are extremely difficult to unlearn.

Employing these theoretical ideas, a stereotypical association (such as "Black" and "aggressiveness") might be stored in semantic memory and automatically activated, producing an implicit stereotype effect. This was demonstrated by Devine (1989). White participants were asked to generate the features of the Black stereotype, and also to complete a prejudice questionnaire. Devine found that both the low- and high-prejudiced individuals knew the characteristics of the Black stereotype. In the next phase of the study the participants rated the hostility of a person only referred to as Donald, described in a 12-sentence paragraph as performing ambiguously hostile behaviours such as demanding his money back on something he had just bought in a store. Before the description, words related to the Black stereotype were rapidly displayed on the screen but too briefly to be consciously recognized. This automatic activation of the stereotype was shown to affect the judgement of Donald's hostility by both the low- and high-prejudiced participants. Finally, the participants were asked to anonymously list their own views of Black people. Low-prejudice individuals gave more positive statements and more beliefs (such as "all people are equal") than traits, whereas high-prejudice participants listed more negative statements and more traits (such as "aggressive").

Devine explained these results by arguing that, during socialization, members of a culture learn the beliefs existing in that culture concerning different social groups. Owing to their frequency of occurrence, stereotypical associations about people from the stereotyped group become firmly-established in memory. Owing to their widespread existence in society, more-or-less everyone in the culture, even the non-prejudiced individual, has the implicit stereotypical associations available in semantic memory. Consequently, the stereotype is automatically activated in the presence of a member of the stereotyped group, and has the potential to influence the perceiver's thought and behaviour. However, people whose personal beliefs reject prejudice and discrimination may seek to consciously inhibit the effect of the stereotype in their thoughts and behaviour. Unfortunately, as described above, conscious processing requires the allocation of attentional resources and so the influence of an automatically activated stereotype may only be inhibited if the person is both aware of its potential bias on activation and is motivated to allocate the time and effort to suppress it and replace it in their decision-making with an intentional non-stereotypical judgement. Devine (1989: 15) viewed the process of asserting conscious control as "the breaking of a bad habit".

It has been argued that conscious attentional resources are only employed when necessary, with the perceiver acting as a "cognitive miser" (Fiske and Taylor, 1991): as a result, Macrae et al. (1994) argued that stereotypes could be viewed as efficient processing "tools", avoiding the need to "expend" valuable conscious processing resources. Yet, Devine and Monteith (1999) argued that they can be consciously suppressed when a non-prejudiced perception is sought. Also an implicit stereotype is only automatically activated when the group member is perceived in terms of a particular social meaning (Macrae et al.,

1997) so automatic activation is not guaranteed on presentation of a group member (Devine and Sharp, 2009). Devine and Sharp (2009) argued that conscious and automatic activation are not mutually exclusive but in social perception there is an interplay between the two processes. Social context can also influence automatic activation so that, in the context of "prisoners" there is a Black stereotype bias (compared with White) but not in the context of "lawyers" (Wittenbrink et al., 2001). Indeed, Devine and Sharp (2009) argued that a range of situational factors and individual differences can affect automatic stereotype activation, and conscious control can suppress their effects on social perception. However, Bargh (1999) was less optimistic than Devine in the ability of individual conscious control to suppress automatically activated stereotypes, and proposed that the only way to stop implicit stereotype influence was "through the eradication of the cultural stereotype itself" (Bargh (1999: 378). Rather than the cognitive miser model of cognitive processing, Bargh proposed the "cognitive monster", arguing that we do not have the degree of conscious control, which Devine proposes, to mitigate the influence of implicit stereotypes (Bargh and Williams, 2006; Bargh, 2011).

Greenwald and Banaji (1995) called for the greater use of indirect measures of implicit cognition to demonstrate the effect of activation outside of the conscious control of the perceiver. They were particularly concerned about implicit stereotypes, arguing that the "automatic operation of stereotypes provides the basis for implicit stereotyping", citing research such as that of Gaertner and McLaughlin (1983). In this latter study, despite participants scoring low on a direct self-report measure of prejudice, they still reliably reacted quicker to an association between "White" and positive attributes, such as "smart", compared with the pairing of "Black" with the same positive attributes. Thus, they concluded that the indirect reaction time measure was identifying an implicit stereotype effect. Consequently, Greenwald et al. (1998) developed the Implicit Association Test (or IAT). This word-association reaction time test presents pairs of words in a sequence of trials over five stages, with each stage examining the reaction time to different combinations of word pairings. From the results at the different stages, the reaction time to various word associations can be examined. For example, the poles of the age concept, "young" and "old", can be sequentially paired with "good" and "bad" to see if the reaction times to the young-good and/or the old-bad pairing are reliably faster than alternative pairings indicating evidence of the implicit stereotype of age. As a technique the IAT can be applied to any word pair combination and as a result can be used to examine a range of implicit stereotypes, such as "White" and "Black" for ethnic stereotyping, or "men" and "women" for gender stereotyping, paired with any words associated with stereotypical attributes, such as aggression or dependence. The results have been quite dramatic. The subsequent use of the IAT has consistently demonstrated implicit stereotyping for a range of different social categories, particularly gender and ethnicity (Greenwald et al., 2015). Implicit stereotyping is now viewed as one aspect of implicit social cognition that is involved in a range of social judgements (Payne and Gawronski, 2010).

Criticisms of the findings of the IAT have questioned whether it is actually identifying a specific unconscious prejudice, unrelated to conscious judgement (Oswald et al., 2013) or, as Devine (1989) suggested, simply knowledge of a cultural association that may be controllable and inhibited in decision-making (Payne and Gawronski, 2010). In support of the IAT, Greenwald et al.'s (2009) meta-analysis of 184 IAT studies showed that there was predictive validity of the implicit associations to behavioural outcomes across a range of subject-areas, and Greenwald et al. (2015) claim this can have significant

societal effects. As a consequence, if implicit stereotyping indicates a potentially-uncontrollable cognitive bias, the question then arises as to how to deal with the outcomes of it in decision-making, particularly for a person genuinely striving for a non-prejudiced judgement. Overt prejudice has been tackled by a range of socio-political measures from anti-discrimination laws to employment interviewer training, but interventions essentially seek to persuade or compel individuals to consciously act in a non-prejudiced way. Lai et al. (2016) examined a range of intervention techniques to reduce implicit racial prejudice, such as exposure to counter-stereotypical exemplars or priming multiculturalism, but the conclusions were somewhat pessimistic. Different interventions had different effects on the implicit stereotype (as measured by the IAT). For example, a vivid counter-stereotypical example (which the participants read)—imagining walking alone at night and being violently assaulted by a White man and rescued by a Black man—was quite effective. However, of the nine interventions examined by Lai et al. (2016), all were effective to some extent but subsequent testing showed that the beneficial effect disappeared within a day or so. The authors concluded that, while implicit associations were malleable in the short term, these (brief) interventions had no long term effect. This could indicate that implicit stereotypes are firmly established and may only be responsive to intensive and long-term interventions (Devine et al., 2012). Lai et al. (2016) also suggest that children may be more susceptible to implicit stereotype change than adults.

The problem is that if people are not consciously able to change their implicit "bias", to what extent are they responsible for actions based on these implicit stereotypes? Law Professor Krieger (1995) argued that lawmakers and lawyers should take account of psychological explanations of implicit bias in their judgements. For example, in a study by Cameron et al. (2010) participants rated the responsibility of a White employer who sometimes discriminated against African Americans, despite a conscious desire to be fair. When this discrimination was presented as resulting from an unconscious bias, that the employer was unaware of, then the personal responsibility for the discrimination was viewed as lower by the participants. However, being told that the implicit bias was an automatic "gut feeling" that the employer was aware of, but found difficult to control, did not produce the same reduction in moral responsibility. This also has potential legal significance (Krieger and Fiske, 2006), as the law has traditionally assumed that a discriminatory act is the responsibility of the individual under-taking that act, with the assumption of an underlying discrimi-natory motivation (an intention). The effect of an implicit stereotype bias may be a discriminatory action that the individual neither intended nor was conscious of.

Implicit stereotype bias provides a challenge to the individual as the sole source and cause of their thoughts and actions. In a huge study of over two hundred thousand participants, all citizens of the USA, Axt et al. (2014) employed the MC-IAT, a variant of the IAT, to examine implicit bias in the judgement of ethnic, religious and age groups. Whilst participants showed in-group favouritism, consistent hierarchies of the social groups emerged in their response times. For ethnicity, in terms of positivity of evaluation, Whites were highest, followed by Asians, Blacks and Hispanics, with the same order obtained from participants from each of the ethnic groups. For religion, a consistent order of Christianity, Judaism, Hinduism and Islam was produced. For the age study, positive evaluations were associated with youth, with a consistent order of children, young adults, middle-aged adults, and old adults, across participants of all ages, from their teens to their sixties. Axt et al. argued that the consistent implicit evaluations reflect cultural hierarchies of social power (and social

structures) "pervasively embedded in social minds" (Axt et al., 2014: 1812). They also suggest that these implicit biases might "not be endorsed and may even be contrary to conscious beliefs and values" (Axt et al., 2014: 1812). The focus on cognitive bias, with its implication of an individual's biased judgement has tended to ignore the importance of culture in cognition. It is this issue that is now considered here.

Implicit cognitive "bias"
Implicit stereotypes are referred to in the literature, and taught to psychology students, as a cognitive bias (Fiske and Taylor, 2013). When, in the past, only a specific group of people were assumed to stereotype (such as authoritarians or the cognitively simple) then they could be viewed as biased in terms of the liberal views of the rest of the population. However, as Fiske and Taylor (2013) claim that now only 10% of the population use overt stereotypes in liberal Western democracies, the major issue is the implicit stereotypes that could affect us all. Indeed, some psychologists (who the reader rightly infers to be supporters of egalitarian values) are willing to reveal examples of their inadvertent use of implicit stereotypes in their own lives—to their chagrin (for example, Stainton Rogers, 2003: 301). Now the assumption is that implicit stereotypes can affect everyone. This makes the use of the term cognitive "bias" problematic when it is universally applied, particularly as it contains the implication of an unconscious cognitive "failing" of the individual (a "cognitive monster" within them), especially given the unsuccessful attempts to correct it, noted above. There also arises the question of how an unbiased judgement can be defined. This idea of an implicit stereotype as a cognitive bias is challenged here.

A wheel is said to be biased if it wobbles on an axle (when others do not). Adjusting it or correcting the imperfections makes it "true" and it is able to run smoothly and straight on the axle. Indeed, the word bias derives from the word "oblique" (for a diagonal thread in weaving) or deviating from the perpendicular. In human social terms, the idiomatic "straight (or strait) and narrow" view might be based on "self-evident truths" (to quote the Declaration of Independence of the USA) rooted in religious or philosophical beliefs, which essentially provide a position from which all other views are biased. Yet, unlike "true" wheels and "fair" coins, there is not an absolute moral standard that is universally accepted, with a long philosophical debate ranging from Plato and Kant to Hume about the issue. Different cultures —as nation states—have different belief systems that are conventionalised into different national legal systems, with dynamically changing laws. Despite the United States Constitu-tion, there are many differences between the views of the Republican and Democratic Parties and their conservative and liberal supporters, and there is a constant political interplay between them about what, in terms of another idiom, is "good and proper" thinking. Recently, the psychologist Haidt (2012) has examined the difference between liberals and conservatives in the USA in terms of their moral foundations. Conventional wisdom is also about both power and politics and in modern times has also been challenged (and changed) by social movements, such as civil rights and women's liberation. Thus, in human terms a "biased" view is often one that differs from the agreed position of a powerful group in a society, with power relations often considered in the sociology of stereotyping (for example, Pickering, 2001), but much less so in the cognitive research. In many cases throughout history, dissenters (such as heretics or dissidents) have been severely punished, imprisoned and put into "psychiatric" institutions, for their unconventional "biased" views.

Furthermore, not all implicit stereotypes have the same cultural value. Consider the associations of "artists" with "creativity" and

"women" with "dependence". Both associations are overgeneralisations and can be labelled as stereotypes. In this sense they are both cognitive "biases". Yet there is no large body of psychological research challenging the stereotype of the creative artist. This is because the two associations differ significantly in their socio-cultural and political meaning. The latter presents a representation of women (common in the past) which is no longer acceptable in a modern liberal democracy where generations of women have politically fought hard to overcome discrimination and achieve equality. Not surprisingly, the majority of the research into stereotyping in the psychological literature has focused on very specific topics: ethnicity or race, gender, sexuality, disability and age. These are all critical issues in the political debates during the last century in Western societies, particularly the USA. Conventional views about these social groups have also undertaken radical change in line with the greater concerns about reducing discrimination and promoting equality. As a result the common views (and associated descriptive terminology) of only a past generation or two are now socially unacceptable and often illegal. It is not unusual to hear modern egalitarian adults discuss with horror the racist or homophobic views they heard at the feet of their grandparents' generation. These topics continue to be of significance in an ongoing political discussion about anti-discrimination and equality in modern Western democracies.

Finally, human cognitive abilities have evolved for a purpose, and implicit associations guiding rapid decision-making have a survival benefit. Fox (1992) argued that this form of pre-judgement (rather than culturally based intergroup prejudices) has evolutionary value. Learning an association of large animals with danger might be "biased" against harmless large animals (who we run away from needlessly) but that is a very small cost to pay compared to a life-saving rapid decision to get out of the way of a dangerous beast. Indeed, Todd *et al.* (2012) argued that it is our ability to make "fast and frugal" strategic (heuristic) judgements that make humans smart. Making decisions using simple associations, based on factors such as recognition or familiarity, may not always result in a logically "correct" answer but can be highly successful heuristics, as research in topics such as economics and investment decision-making, emergency medicine and consumer behaviour have all shown (Gigerenzer and Gaissmaier, 2011). The model of the person emerging from the implicit stereotyping research appears to characterise the fair-minded individual as wrestling with an implicitly biased cognitive monster within them. However, it is argued here that this is a false image. We learn the cultural mores of our society through socialisation and daily communication with other members of the culture. We may not approve of all aspects of our culture (and indeed might strongly object to some) but cultural knowledge—just like other knowledge—is crucial to our pragmatic functioning in society. The wide range of semantic associations we learn in our culture can successfully guide our judgements from what to wear at a job interview, which side of the road to drive on, and how to talk to the boss. In order to change the specific set of implicit associations which we find consciously objectionable, it may be better to explore ways of changing the culture to undermine these specific associations, rather than focusing on the inferred "bias" of human cognition: as is argued from the "predictive brain" model below, human cognition is functionally driven to pick up regularities and develop implicit associations from the world around us.

The predictive brain

It has been proposed that human brains are "prediction machines" (Clark, 2013: 181), in that experience develops expectations. Perception operates by employing prior probabilities that are efficiently deployed to reduce the processing requirements of treating each new experience as completely new. While explored mostly with basic object perception, Clark (2013) argued that it is applicable to social perception, and Otten *et al.* (2017) have applied it to social knowledge. For Clark (2014) perceiving is predicting. For example, we are able to quickly and efficiently recognize a friend we have arranged to meet outside a restaurant, even from quite a distance. Through repeated experience of the friend we have developed a sophisticated prediction based on a range of cues from their gait to their favourite coat. Usually, this prediction is correct and it is the person we expected. The dynamic of the predictive brain is to minimise the error of the prediction, that is, the difference between the prediction and the experienced event. Every now and again we are "surprised"—we mistake a stranger for the friend—and this instance of "surprisal" (an engineering term for the error) will also have an incremental effect on the probabilities (and we might be a little more careful when we next meet the friend). The brain seeks to minimise "surprisal" by a constant process of updating probabilities with each experience. However, an occasional error—as only one instance—will normally only have a small effect on the prior probabilities that have been developed over multiple successful perceptions. In this model of the brain, cognitive bias is not an inaccurate deviation from a "true" position, but an expectation or prediction based on the prior probabilities that have developed through experience. Prediction is not about being correct every time—but is about minimising error and maximising predictive accuracy. This process follows Bayes' Theorem, which expresses a probability of one event (A) given that another event (B) has occurred (such as it being the friend, given the familiarity of the coat and hairstyle observed). This is referred as "likelihood". Human perception operating according to Bayesian decision-making has been studied in both psychology and economics, so the predictive brain model is also referred to as the "Bayesian brain" (El-Gamal and Grether, 1995; Bubic *et al.*, 2010). The implicit semantic associations of "bread" and "butter" or "table" and "chair" (Neely, 1977) have developed through their repeated co-occurrence during our experience of the world. Clearly in ancient Japan (without bread and butter or Western-style tables and chairs) these specific implicit associations did not develop. In social perception we can ask: what is the probability of this man being a basketball player given that he is a tall, Black professional sportsman? This likelihood is based on prior probabilities—which come from experience or knowledge of the culture—so the likelihood could be judged differently by a person from the USA compared with a person from Kenya.

Allport (1979: 191) proposed that stereotypes were "exaggerated beliefs" associated with a social group, citing "all lawyers are crooked" as an example. The idea that stereotypes involve a belief that *all* members of the category share an attribute has persisted in the cognitive research (Hinton, 2000). However, Allport (1979: 189) also stated that a stereotype is "a generalized judgement based on a certain probability that an object of a class will possess a given attribute". This is not the same. The assumption that stereotypes involve "all" judgements presents them as rigid and fixed, yet the probabilistic association of a stereotyped group member and a specific attribute does not. The presence of an honest lawyer demonstrably *proves* the former "all" statement to be an incorrect generalization. In the latter case, which follows from the predictive brain model, the experience of an honest lawyer will only adjust the probabilities according to Bayes' theorem, making it slightly less likely that the next (unknown) lawyer will be predicted to be crooked.

In a well-known study Kahneman and Tversky (1973: 241) gave participants a description of Jack that matched the stereotype of an engineer:

> Jack is a 45-year-old man. He is married and has four children. He is generally conservative, careful, and ambitious. He shows no interest in political and social issues and spends most of his free time on his many hobbies which include home carpentry, sailing and mathematical puzzles.

They were then asked to predict the probability of Jack actually being an engineer in a room of 30 engineers and 70 lawyers. Participants tended to ignore the base-rate probabilities (0.3 for engineer and 0.7 for lawyer) but made their judgements on the stereotypicality of the description. Kahneman and Tversky argued that the participants were making their judgements on the similarity of the description to the engineer stereotype, which they called "the representativeness heuristic", and not on the base-rate probabilities. They argued that this strategy was not as good as using the base-rate probabilities as the description may not be valid and furthermore it could match more of the lawyer group as there are simply more of them. However, they admitted that a Bayesian prediction could produce the likelihood of Jack being an engineer if the description was accurate and diagnostic. This highlights a key problem of arguing that people's judgements are "biased" compared to an "accurate" measurement. Outside the psychological laboratory people almost never know the base-rate probabilities (Todd *et al.*, 2012) and learnt associations are often all they have to go on. In an attempt to find accurate demographic information about engineers, I discovered that 80% of engineering students in the USA are men (Crawford, 2012)—which is not diagnostic in this case—but could find no data on the overall proportion of engineers who are uninterested in politics or enjoy mathematical puzzles. In many cases like this, accurate demographic data is unavailable, either because it is not there or because we do not have the time and motivation to find it—we can only rely on our general knowledge of engineers. The Bayesian brain develops its statistical probabilities from experience of engineers—such as the engineers encountered in life and learnt about through the media. The likelihood that an engineer is a man who is uninterested in politics and likes mathematical puzzles does not mean that *all* engineers must have these attributes, simply that these are frequently encountered in engineers in the social world, such as the engineer Howard Wolowitz in the popular US sitcom *The Big Bang Theory*, 2007. Thus, the predictive brain, operating through past experience and subtly adapting to each new experience, is a pragmatically functional system rather than being "biased" by an all-or-none overgeneralization.

Consider the following example where I could find some demographic information[1]. There are 70 professional golfers and 30 professional basketball players in a room (all men and from the USA). The only available information is that Tom is 193 cm tall (6′4″). What is the probability that he is a basketball player? From Kahneman and Tversky (1973), we can infer that a participant will respond, using the representativeness heuristic, that Tom is probably a basketball player on the learnt association that "basketball players are tall". Using only the base-rate probabilities Tom should be predicted to be a golfer. However, a Bayesian analysis of the demographic data agrees with the representativeness heuristic that it is very likely that Tom is a basketball player. Rather than assuming that human cognition is statistically naive, an alternative explanation is that people are unconsciously Bayesian and they normally assume that a description identifying learnt implicit associations is accurate

and diagnostic (unless they consciously decide otherwise). Kahneman (2011: 151) acknowledges the link between height and basketball players as an example of where representativeness can lead to a more accurate than chance guess of an athlete's sport. Outside of the psychological laboratory it may be that a limited description is all the information people have to go on. Indeed, Jussim (2012) argues that when a perceiver has almost no information about a person except, say, a social category ("This person is an engineer") then they may employ stereotypical associations, based on social knowledge, to make predictions about them ("Engineers are not interested in politics"), which may well be accurate. However, in an encounter with the specific person, they will learn new information to adjust this view if the prediction is not supported.

Jussim (2012: 159) argued, in agreement with Kelly (1955), that people operate as naïve scientists, seeking to make accurate predictions of people and events based on expectation and, in the research focus on bias, the evidence that social perception is generally accurate has been ignored, with various independent factors often conflated in the discussion of stereotype accuracy. For example, if a perceiver Ben predicts, on the stereotypical association of a social group and underachievement, that Joe (a member of the group) will not get into the top university he has applied for, and Joe is rejected by the university, then Ben's social perception is accurate. However, this does not relate to Ben's belief about why Joe wasn't admitted or the actual reason why Joe was not admitted. Ben could be prejudiced against the social group (believing the stereotype) but, alternatively, he might be a fair-minded person who believes that the university is prejudiced against the group in its procedures. Also the university might have rejected Ben either as it is prejudiced in its selection or, alternatively, has a fair-assessment system and Ben is rejected for reasons unrelated to his group membership. These additional factors do not mitigate the evidence that Ben's social judgement was correct. Jussim (2012: 155) challenged the researchers who criticize the "permissibility" of relying on stereotypes in judgement social judgement—arguing for a moral imperative that stereotypes should not be employed in social judgements—in their rejection of the accuracy data.

A key point to note here is that the predictive brain operates on the state of the world *as it is experienced* and not on the state of the world *as we believe it should be*. Working towards gender equality and encouraging more women into engineering is a key aim in many Western societies, but that admirable social and political goal should not lead us to misunderstand the unconscious working of the predictive brain. Indeed, according to Crawford's (2012) figures, the probabilistic association of "engineer" and "man" is an accurate reflection of the "true" state of the USA in 2012 where 80% of the recruits to the profession are men. A second important point is that the Bayesian brain seeks predictive validity through the picking up of regularities (to form associations) on the basis of experience. Diversity, or counter-stereotypical examples (such as encountering a woman engineer) will reduce the probability of an association (between "engineer" and "man"), but only to the degree that they are experienced. Whereas the presence of even a single female engineer disproves the assertion that "all engineers are men"— and demonstrates that gender is not a relevant factor in engineering ability—the presence of only one female engineer (where all the rest are men) will only have a small effect on the predictive probability of an engineer being a man. The implication from the predictive brain model is that when there are more women engineers, who then become more visible in everyday life (and in the media) then the implicit stereotypical association of "engineer" and "man" will change (Weber and Crocker, 1983).

The predictive brain, as a perceptual mechanism, is directed solely by the minimization of surprise. It does not make a moral judgement or provide an explanation for the state of the world. It simply seeks to make accurate predictions. In a study on language learning, Perfors and Navarro (2014) argued that the Bayesian brain learns through a process of iterative learning (from other members of the community). Whereas previous researchers have argued that it is solely the structure of language that structures the meanings acquired, Perfors and Navarro (2014) argued that the structure of the external world (and the meanings within it) will also influence the process. We don't simply learn that an engineer, by definition, designs and builds systems but also that, in the external world, they are mostly men. Thus, semantic knowledge acquired will be shaped by the meaning structure communicated. As long as the things people talk about reflect the relationships of those things in the external world then the semantic relationships learnt will reflect the meanings present in the external world. Thus, knowledge of the relationship between concepts will be acquired from the meanings communicated by others. Furthermore, the proposal of a Bayesian brain does not require that it operates in an optimal (or rational) manner—simply that a Bayesian model best represents its behaviour (Tauber et al., 2017). Learning for the Bayesian brain involves testing predictions (hypotheses) by using the data obtained from the world and applying Bayes' theorem to develop probabilities (Perfors, 2016). For the predictive brain, the degree to which implicit stereotypes are learnt and employed depends on the probabilities with which the implicit associations between the social category and an attribute are expected and experienced in communication. It is this world of the social perceiver that is considered now.

Implicit stereotypes and "culture in mind"

Implicit stereotypes, like other implicit associations can be viewed as cultural knowledge or folk wisdom that the person acquires through their experience in a culture (Bruner, 1990). The idea that stereotypical associations are cultural in origin was proposed in the early work on stereotypes, but has tended to be ignored in the focus on the fallacy or bias of individual cognition. Journalist and political commentator Walter Lippmann is usually seen as stimulating the academic study of stereotyping with his 1922 book *Public Opinion* (Hinton, 2000). While Lippmann used the term "stereotype" familiar to him from newspaper printing, he saw it as a cultural phenomenon: "we tend to perceive that which has been picked out in the form *stereotyped for us by our culture*." (Lippmann, 1922: 81; my italics) In Lippmann's view it is the culture that is creating the stereotype, not the individual (Hinton, 2016). As Allport (1979: 189) pointed out: stereotypes "manifestly come from *somewhere*". To illustrate this, we can examine the origin of the associations identified in the Princeton studies, discussed at the beginning of this article, by considering the example of the English. As Hinton (2016) has argued, the selected attributes reflect the notion of the English gentleman, a common representation of the Englishman in the American media of the first half of the twentieth century, and hence familiar to the exclusively male, upper-class Princeton student participants who, if they had encountered English people it is likely that they would be from the same class demographic as themselves. It is also likely that these participants did not consider (nor were they asked to do so) a range of categories of English people, such as women or the working classes, so, not surprisingly, tended to focus on the specific and familiar representation of the English defined for them by their culture (to paraphrase Lippmann). By 1969, the image of the English gentleman had become rather archaic and even a figure of fun in both the British and American media

(Hinton, 2016) and the selected English attributes had changed. Also, a crucial point to note is that the student participants were only asked "to select those [attributes] which seem to you to be typical" of the group (Katz and Braly, 1933: 282). Even so, some students refused to do the task in 1951 and 1969 (Brown et al., 1987), which indicates that, even for the students who had agreed to take part in the study, there was no evidence that the selected attributes represented their own personal attitudes, thus the responses did not reflect a fallacy or a cognitive bias of the participants. To perform the task with no information except the category name, the students may have simply drawn on attributes they knew to be commonly circulating about the English in their culture. The most popular attribute in 1933 for the English was "sportsmanlike", and this might even have shown up in the IAT if it had been available at the time. Yet this does not mean that the students viewed all English people as sportsmanlike. However, the sportsmanlike English gentleman was a familiar trope in American popular culture at the time, typified by actor Ronald Colman in Hollywood movies such as *The Dark Angel*, 1925, and *Bulldog Drummond*, 1929. By 1969, "sportsmanlike" had dropped out of the Princeton top five attributes for the English (Karlins et al., 1969). We can take Allport's example of the "crooked lawyer" stereotype as a second example. A person with no personal antipathy to lawyers, and well-aware that they are a highly regulated profession of mostly honest people, might make the prediction that when a lawyer character appears in a popular crime drama that they will (probably) be crooked from the experience of lawyers in famous movies such as *The Godfather* series, 1972–1990, and television programs such as *Breaking Bad*, 2008–2013, (along with the spin-off series about a crooked lawyer, *Better Call Saul*, 2015).

As Devine (1989) has argued, well-learnt associations picked up during socialization form implicit stereotypes even for the individual seeking non-prejudiced views. It is argued here that the predictive brain model provides the mechanism for this. The process of picking up associations probabilistically is happening unconsciously through Bayesian principles throughout a person's life within a culture. Yet culture is neither monolithic nor fixed and unchanging. People are active in the construction both of their social world and their media environment (Livingstone, 2013; Burr, 2015). As Smith (2008: 51) points out "In reality, people's social environments are probably best characterized as *social networks*. People have links of acquaintanceship, friendship, etc. to particular other people, which interconnect them in a complex web". Within any society, there will be different social networks of this kind communicating different social representations about social groups. According to Moscovici (1998), it is these shared representations that define a culture or subcultural group. Different cultural groups will differ ideologically through their position in society and the representations that circulate in the communication within their social network. While one cultural group may be actively promoting one representation (such as "immigrants" are "a great economic benefit to our society and add to the diversity of our culture") through a range of communications, such as television, newspaper and social media, another group may be promoting an alternative representation (such as "immigrants" are "a burden on society, taking jobs and undermining our culture"). In the communication within any social network there will be regular and consistent associations between social groups and attributes, which will be picked up by it members, through the working of the predictive brain. The extent to which individuals share implicit associations will depend on the hegemonic social representations within the society across cultural groups (Gillespie, 2008), such as a positive belief in democracy and a negative view of communism, which are prevalent in the wider social institutions within a nation, and

examined in the sociological study of stereotypes (for example, Pickering, 2001).

The role of stereotypes in communication within a social network was demonstrated by Kashima and colleagues (Kashima and Yeung, 2010; Kashima *et al.*, 2013) in their research on the serial retelling of stories. The results showed that stereotype-consistent information was emphasized. Even though stereotype-inconsistent information attracted attention it was not necessarily passed on. Thus, the story became more stereotypical and consistent in the serial retelling. They argued that "stereotypes can be thought of a significant cultural resources that help us to transmit cultural information" (Kashima and Yeung, 2010). Within a social network common understandings are developed via the use of stereotypes. Members of the culture assume a knowledge of the stereotype in other group members, which facilitates social interaction, but potentially also helps to maintain the stereotype, even in the face of inconsistent information. From this research, it can be argued that the analysis of implicit stereotypes should focus on the communication of meaning within a social network, rather than considering them as a "bias". The complex dynamics of the individual within a social network (for example, Christakis and Fowler, 2009) needs to be considered in investigating the formation, transmission and maintenance of implicit stereotypes.

In the modern world of the twenty-first century, the options available for people to construct their social environments have radically increased (Giddens, 1991). The media has rapidly expanded through multiple television channels, a proliferation of media outlets, and the development of social media via the internet. While this offers the potential for people to engage with a diversity of representation and counter-stereotypical information, it also allows people to remain in an ideological subculture, communicating with like-minded people where specific representations of cultural others are constantly being circulated unchallenged within the social network. In terms of the predictive brain, implicit associations will develop from the consistent messages people receive in their everyday lives. If certain implicit stereotypes are deemed unacceptable then it will only be when people experience consistent counter-stereotypical information over a long period of time that these associations will be probabilistically undermined. For this to be achieved, everyday experience has to involve necessarily (but not sufficiently) exposure to alternative representations and counter-evidence to these specific implicit stereotypes, rather than people only experiencing the consistent representations about social groups circulating within a particular culture, social network or social media "bubble".

Conclusion

Over the last 30 years stereotype research has focused on implicit stereotypes, particularly using the IAT, which have been interpreted as revealing an implicit or unconscious cognitive bias, even for the consciously fair-minded person. Despite research questioning the predictive validity of the IAT as a method of revealing unconscious prejudice (for example, Oswald *et al.*, 2013), the focus of implicit stereotypes has dominated the psychology of stereotyping in the twenty-first century (Fiske and Taylor, 2013). However, it is argued here that implicit stereotypes, as attributes associated with social groups, do not indicate an unconscious cognitive "bias" (a "cognitive monster") within the fair-minded person but are learnt associations arising from the normal working of the predictive brain in everyday life. These associations are based on information circulating within the person's culture, and the associations are probabilistically detected by the predictive brain: as such they can be characterised

as "culture in mind" rather than an individual bias. According to the predictive brain model, when the culture changes then the implicit stereotypes of its members will change (albeit slowly for some associations). Therefore, to properly understand the nature of implicit stereotypes, the cognitive research needs to be combined with the study of the dynamics of culture, to understand the specific associations prevalent in the communication within a culture and their implicit influence on the members of that culture.

Notes

1 I worked with the following data. American golfers are of normal height (say, mean of 173 cm or 5 feet and 10 inches), professional basketball players in the USA are tall (say, a mean of 200 cm of 6 feet 7 inches), assumed standard deviations of 10 cm or 4 inches for both groups, and heights to be normally distributed.

References

Allport GW (1979[1954]) *The Nature of Prejudice*. Addison-Wesley: Reading, MA.

Adorno TW, Frenkel-Brunswik E, Levinson DJ and Sanford RN (1950) *The Authoritarian Personality*. Harper and Row: New York.

Axt JR, Ebersole CR and Nosek BA (2014) The rules of implicit evaluation by race, religion, and age. *Psychological Science*; **25** (9): 1804–1815.

Bargh JA (2011) Unconscious thought theory and its discontents: A critique of the critiques. *Social Cognition*; **29** (6): 629–647.

Bargh JA (1999) The cognitive monster: The case against controllability of automatic stereotype effects. In: Chaiken S and Trope Y (eds). *Dual Process Theories in Social Psychology*. Guilford: New York, pp 361–382.

Bargh JA and Williams EL (2006) The automaticity of social life. *Current Directions in Psychological Science*; **15** (1): 1–4.

Brown R (1965) *Social Psychology*. Collier-Macmillan: London.

Brown BL, Williams RN, Norton RF and Barrus GS (1987) Semantic space comparisons of cross-cultural stereotyping. *Deseret Language and Linguistic Society Symposium*; **13** (1): Article 3.

Bruner JS (1990) *Acts of Meaning*. Harvard University Press: Cambridge, MA.

Bubic A, von Cramon DY and Schubotz RI (2010) Prediction, cognition and the brain. *Frontiers in Human Neuroscience*; **4** (25): 1–15.

Burr V (2015) *Social Constructionism*, 3rd edn, Routledge: London.

Cameron CD, Payne BK and Knobe J (2010) Do theories of implict race bias change moral judgements? *Social Justice Research*; **23** (4): 272–289.

Christakis NA and Fowler JH (2009) *Connected: The Surprising Power of our Social Networks and how they Shape Our Lives*. Little & Brown: New York.

Clark A (2014) Perceiving as predicting. In: Stokes D, Matthen M and Biggs S (eds). *Perception and Its Modalities*. Oxford University Press: New York, pp 23–43.

Clark A (2013) Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*; **36** (3): 181–204.

Crawford M (2012) Engineering still needs more women, https://www.asme.org/career-education/articles/undergraduate-students/engineering-still-needs-more-women, accessed 12 December 2016.

Collins AM and Loftus EF (1975) A spreading activation theory of semantic processing. *Psychological Review*; **82** (6): 407–428.

De Deyne S, Navarro DJ, Perfors A and Storms G (2016) Structure at every scale: A semantic network account for the similarities between unrelated concepts. *Journal of Experimental Psychology: General*; **145** (9): 1228–1254.

Devine PG (1989) Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality & Social Psychology.*; **56** (1): 5–18.

Devine PG, Forscher PS, Austin AJ and Cox WTL (2012) Long-term reduction in implicit race bias: A prejudice habit-breaking intervention. *Journal of Experimental Social Psychology*; **48** (6): 1267–1278.

Devine PG and Monteith MJ (1999) Automaticity and control in stereotyping. In: Chaiken S and Trope Y (eds). *Dual-process Theories in Social Psychology*. Guilford: New York, pp 339–360.

Devine PG and Sharp LB (2009) Automaticity and control in stereotyping and prejudice. In: Nelson TD (ed). *Handbook of Prejudice, Sereotyping, and Discrimination*. Taylor and Francis: New York, pp 61–88.

El-Gamal MA and Grether DM (1995) Are people Bayesian? Uncovering behavioral strategies. *Journal of the American Statistical Society*; **90** (432): 1137–1145.

Fiske ST and Taylor SE (2013) *Social Cognition: From Brains to Culture*, 2nd edn, Sage: London.

Fiske ST and Taylor SE (1991) *Social Cognition*. 2nd edn, McGraw-Hill: New York.

Fox R (1992) Prejudice and the unfinished mind: A new look at an old failing. *Psychological Inquiry*; **3** (2): 137–152.

Gaertner SL and McLaughlin JP (1983) Racial stereotypes: Associations and ascriptions of positive and negative characteristics. *Social Psychology Quarterly*; **46** (1): 23–30.

Ghosh VE and Gilboa A (2014) What is a memory schema? A historical perspective on current neuroscience literature. *Neuropsychologia*; **53**, 104–114.

Giddens A (1991) *Modernity and Self-identity: Self and Society in the Late Modern Age*. Polity Press: Cambridge, UK.

Gigerenzer G and Gaissmaier W (2011) Heuristic decision making. *Annual Review of Psychology*; **62**, 451–482.

Gilbert GM (1951) Stereotype persistence and change among college students. *Journal of Abnormal and Social Psychology*; **46** (2): 245–254.

Gillespie A (2008) Social representations, alternative representations and semantic barriers. *Journal for the Theory of Social Behaviour*; **38** (4): 375–391.

Greenwald AG and Banaji MR (1995) Implicit social cognition: Attitude, self-esteem, and stereotypes. *Psychological Review*; **102** (1): 4–27.

Greenwald AG, Banaji MR and Nosek BA (2015) Statistically small effects of the Implicit Association Tst can have socially large effects. *Journal of Personality and Social Psychology*; **108** (4): 553–561.

Greenwald AG, McGhee DE and Schwartz JLK (1998) Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology*; **74** (6): 1464–1480.

Greenwald AG, Poehlman TA, Uhlmann E and Banaji MR (2009) Understanding and using the implicit association test: III. Meta-analysis of predictive validity. *Journal of Personality and Social Psychology*; **97** (1): 17–41.

Haidt J (2012) *The Righteous Mind: Why Good People are divided by Politics and Religion*. Vintage Books: New York.

Hinton PR (2016) *The Perception of People: Integrating Cognition and Culture*. Routledge: London.

Hinton PR (2000) *Stereotypes, Cognition and Culture*. Psychology Press: Hove, UK.

Jussim L (2012) *Social Perception and Social Reality: Why Accuracy Dominates Bias and Self-fulfilling Prophecy*. Oxford University Press: New York.

Kahneman D (2011) *Thinking, Fast and Slow*. Penguin Books: London.

Kahneman D and Tversky A (1973) On the psychology of prediction. *Psychological Review*; **80** (4): 237–251.

Karlins M, Coffman TL and Walters G (1969) On the fading of social stereotypes: Studies in three generations of college students. *Journal of Personality and Social Psychology*; **13** (1): 1–16.

Kashima Y, Lyons A and Clark A (2013) The maintenance of cultural stereotypes in the conversational retelling of narratives. *Asian Journal of Social Psychology*; **16** (1): 60–70.

Kashima Y and Yeung VWL (2010) Serial reproduction: An experimental simulation of cultural dynamics. *Acta Psychologica Sinica*; **42** (1): 56–71.

Katz D and Braly KW (1933) Racial prejudice and racial stereotypes. *Journal of Abnormal and Social Psychology*; **30** (2): 175–193.

Kelly GA (1955) *A Theory of Personality: The Theory of Personal Constructs*. Norton: New York.

Koenig FW and King MB Jr (1964) Cognitive simplicity and out-group stereotyping. *Social Forces*; **42** (3): 324–327.

Krieger LH (1995) The content of our categories: A cognitive bias approach to discrimination and equal employment opportunity. *Stanford Law Review*; **47** (6): 1161–1248.

Krieger LH and Fiske ST (2006) Behavioral realism in employment discrimination law: Implicit bias and disparate treatment. *California Law Review*; **94** (4): 997–1062.

Lai CK, Skinner AL, Cooley E et al (2016) Reducing implicit racial preferences: II. Intervention effectiveness across time. *Journal of Experimental Psychology: General*; **145** (8): 1001–1016.

Lippmann W (1922) *Public Opinion*. Harcourt-Brace: New York.

Livingstone S (2013) The participation paradigm in audience research. *Communication Review*; **16** (1-2): 21–30.

Macrae CN, Bodenhausen GV, Milne AB, Thorn TMJ and Castelli L (1997) On the activation of social stereotypes: The moderating role of processing objectives. *Journal of Experimental Social Psychology*; **33** (5): 471–489.

Macrae CN, Milne AB and Bodenhausen GV (1994) Stereotypes as energy-saving devices: A peek inside the cognitive toolbox. *Journal of Personality and Social Psychology*; **66** (1): 37–47.

Moscovici S (1998) The history and actuality of social representations. In: Flick U (ed). *The Psychology of the Social*. Cambridge University Press: Cambridge, UK. pp 209–247.

Neely JH (1977) Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General*; **106** (3): 226–254.

Nelson TD (2009) *Handbook of Prejudice, Stereotyping, and Discrimination*. Psychology Press: Hove, UK.

Oswald FL, Mitchell G, Blanton H, Jaccard J and Tetlock PE (2013) Predicting ethnic and racial discrimination: A meta-analysis of IAT criterion studies. *Journal of Personality and Social Psychology*; **105** (2): 171–192.

Otten M, Seth AK and Pinto Y (2017) A social Bayesian brain: How social knowledge can shape visual perception. *Brain and Cognition*; **112**, 69–77.

Payne BK and Cameron CD (2013) Implicit social cognition and mental representation. In: Carlston D (ed). *Oxford handbook of social cognition*. Oxford University Press: Oxford.

Payne BK and Gawronski B (2010) A history of implicit social cognition: Where is it coming from? Where is it now? Where is it going? In: Gawronski B and Payne BK (eds). *Handbook of Implicit Social Cognition: Measurement, Theory, and Applications*. Guilford Press: New York, pp 1–15.

Perfors A (2016) Piaget, probability, causality, and contradiction. *Human Development*; **59** (1): 26–33.

Perfors A and Navarro DJ (2014) Language evolution can be shaped by the structure of the world. *Cognitive Science*; **38** (4): 775–793.

Pickering M (2001) *Stereotyping: The Politics of Representation*. Palgrave Macmillan: New York.

Schneider DJ (2004) *The Psychology of Stereotyping*. Guilford: New York.

Shiffrin RM and Schneider W (1977) Controlled and automatic human information processing: II. perceptual learning, automatic attending and a general theory. *Psychological Review*; **84** (2): 127–190.

Smith ER (2008) Rediscovering the emotional aspects of prejudice and intergroup behavior. In: Tropp LR, Finchilescu G and Tredoux C (eds). *Improving Intergroup Relations: Building on the Legacy of Thomas F. Pettigrew*. Wiley: New York, pp 42–54.

Stainton RW (2003) *Social Psychology: Experimental and Critical Approaches*. Open University Press: Maidenhead, UK.

Steele C (2010) *Whistling Vivaldi: and Other Clues to how Stereotypes affect us*. W.W. Norton: New York.

Tajfel H (1969) Cognitive aspects of prejudice. *Journal of Social Issues*; **25** (4): 79–93.

Tauber S, Navarro DJ, Perfors A and Steyvers M (2017) Bayesian models of cognition revisited: Setting optimality aside and letting data drive psychological theory. *Psychological Review*; **124** (4): 410–441.

Todd PM and Gigerenzer Gthe ABC Research Group. (2012) *Ecological Rationality: Intelligence in the World*. Oxford University Press: Oxford.

Weber R and Crocker J (1983) Cognitive processes in the revision of stereotypic beliefs. *Journal of Personality and Social Psychology*; **45** (5): 961–977.

Wittenbrink B, Judd CM and Park B (2001) Spontaneous prejudice in context: Variability in automatically activated attitudes. *Journal of Personality and Social Psychology*; **81** (5): 815–827.

Data availability

Data sharing is not applicable to this article as no datasets were generated or analysed.

Additional information