If Nazi = Red, and Canadian = Red, does Red = Good or Bad? Testing the effects of valenced perceptual cues on Implicit Association Test performance

Michael McCarthy¹

¹ University of the Fraser Valley

Author Note

This research was first presented as a poster at the University of the Fraser Valley's Student Research Day, 2019, and at the North West Cognition And Memory (NOWCAM) conference, 2019. This article is currently a working draft, so any feedback and criticism is highly appreciated. I would like this article to be as accurate as possible, as I hope the ideas presented within can be used to improve implicit social cognition research.

Correspondence concerning this article should be addressed to Michael McCarthy, via E-mail. E-mail: Michael.McCarthy@student.ufv.ca

Abstract

After two decades of research on *implicit social cognition*, it has become clear many of the field's theories and practices need to be redressed. Addressed in the present paper are the assumptions and memetic ideas embedded in the language, theory, and measurement tools of so-called implicit social cognition, their historical and contemporary shortcomings, and solutions to avoid these shortcomings and improve scholarship in this field. Specifically, the present paper recommends researchers and theorists adopt more accurate and epistemically honest language in their work, and determine what measures of implicit social cognition measure through experimental research rather than through empirically unjustifiable assumptions. Contributing to this endeavour, the present paper includes an experiment testing whether performance on the Implicit Association Test (IAT) can be changed through the indirect activation of complex associations. Here it is demonstrated that performance on IAT appears to be unaffected by indirectly activated complex associations, ruling out a possible cognitive mechanism that could contribute to performance on the IAT.

Keywords: Implicit Association Test, Implicit Attitude Change, Attitude

Generalization

Word count: X

If Nazi = Red, and Canadian = Red, does Red = Good or Bad? Testing the effects of valenced perceptual cues on Implicit Association Test performance

Introduction

It is difficult to imagine a psychological construct that has led to as much misunderstanding, misappraisal, misuse, and general confusion, as that of the implicit attitude construct (Greenwald & Banaji, 1995)—particularly because of the magnitude of its popularity (Greenwald & Banaji, 2017; Mitchell & Tetlock, 2017), and the amount of empirical and theoretical scrutiny it has received (e.g., Blanton & Jaccard, 2008; Conrey, Sherman, Gawronski, Hugenberg, & Groom, 2005; Fiedler, Messner, & Bluemke, 2006; Gawronski et al., 2006; Mitchell & Tetlock, 2017; Oswald, Mitchell, Blanton, Jaccard, & Tetlock, 2013, 2015). Yet, more than two decades after its conception, even the basic meaning and operationalization of the term implicit attitude remains unresolved (see Greenwald & Banaji, 2017); with some publications using the term conceptually (implicit attitude = unconscious attitude) (e.g., Johnson et al., 2017), others using the term empirically (implicit attitude = indirectly measured attitude) (e.g., Costantini et al., 2019), and others conflating the outcome of an indirect measurement with an unconscious process (e.g., Malinen & Johnston, 2013). Although disagreements within any field are normal, encouraging even, the lack of consensus on what is meant by the term implicit attitude highlights the carelessness with which it has been used by academics and researchers (see Gawronski et al., 2006; Mitchell & Tetlock, 2017), particularly when the terms *unconscious attitude* and *indirectly measured attitude* have always been available to use in place of the less descriptive, poorly defined, *implicit attitude*. Regretably, such carelessness has led to the implicit attitude construct—and the dual-process typology it is founded on—becoming a meme inside (see Melnikoff & Bargh, 2018; Mitchell & Tetlock, 2017) and outside of academic journals (e.g., Comey, 2015; Gladwell, 2007; Google, 2013; Starbucks, 2018), leading to vast swathes of the population (psychologists included) purporting a

number of empirically unsupported, internally inconsistent, and perhaps untestable, claims about human cognition and behaviour (see Brownstein, Madva, & Gawronski, 2019; Gawronski et al., 2006; Melnikoff & Bargh, 2018; Mitchell & Tetlock, 2017).

Indeed, a number of recent research findings, literature reviews, and meta-analyses, have implicated that there is a growing need for the implicit attitude construct to be reigned in, overhauled, and/or reconceptualized (e.g., Brownstein et al., 2019; Corneille & Stahl, 2019; De Houwer, 2014; Fiedler et al., 2006; Greenwald & Banaji, 2017; Melnikoff & Bargh, 2018; Mitchell & Tetlock, 2017; Oswald et al., 2013, 2015; Payne, Vuletich, & Lundberg, 2017). Although such suggestions have been present since the conception of the implicit attitude construct (e.g., Fazio & Olson, 2003; Karpinski & Hilton, 2001; Kühnen et al., 2001), they have either gone unheeded—as evidenced by continued, empirically unsupported, uses of the term implicit attitude (see Greenwald & Banaji, 2017)—or disregarded for pragmatic reasons (Nosek & Greenwald, 2009). However, in the face of mounting, legitimate, criticisms (e.g., Mitchell & Tetlock, 2017), it is clear that many of the assumptions about the implicit attitude construct need to be addressed, and that some common-ground needs to be reached regarding the operationalization, measurement, and interpretation, of the *implicit* construct and terminology in order to reduce (or eliminate) the misunderstanding, misappraisal, misuse, and general confusion that has been an implicit feature of historical and contemporary implicit cognition scholarship.

The Precisely Imprecise Language of Implicit Cognition

It is difficult to identify the continued motivation to use the word *implicit* in the term implicit cognition, *particularly in a scientific context*, as common usages of the word implicit imply something which may only be revealed by inductive methods. Specifically, something which is implicit is expressed indirectly, thus, it must be discerned through assumption and inference; it is not empirically observable or testable. Thus, to say something is implicit is to express belief in its existence, but it is not to provide evidence of

its existence. Furthermore, because belief is required it is possible for something which is implicit to have multiple simultaneous truths, as different lines of inference may lead to different competing conclusions; all of which must rely on rhetoric, rather than the scientific method. In other words, implicit in the language of implicit cognition is an obfuscation of conciseness, clarity of expression, and meaning. Indeed, this issue is the cornerstone for many of the problems discussed throughout this section.

Of course, explicit reasons behind the decision to choose the word implicit were originally provided in the first paper describing the concept of implicit cognition: Greenwald and Banaji (1995) wrote that their decision to use the term implicit (over other terms such as indirect, unconscious, non-declarative, etc.) was based on their intention to connect research on attitudes, self-esteem, and stereotypes to memory research, which itself was prominently using an implicit-explicit dichotomy to describe different systems of memory (i.e., implicit memory vs. explicit memory). In other words, it was a pragmatic decision; one that also allowed the concept of implicit cognition to cast a wide theoretical net. Specifically Greenwald and Banaji (1995) note that the terms implicit-explicit were used to capture a set of overlapping, dichotomous distinctions mainly used in the psychological literature to describe opposing psychological processes or systems, namely: unaware-aware, unconscious-conscious, intuitive-analytic, indirect-direct, procedural-declarative, and automatic-controlled. Thus, under this framework the concept of implicit cognition was intended to capture a wide range of psychological processes which were unconscious, automatic, and intuitive, which operate outside of awareness, and which can only be measured indirectly; whereas the concept of explicit cognition was intended to capture a wide range of psychological processes which were conscious, controllable, and analytic, which operate within our awareness, and which can be measured directly. This is where the problems begin.

Problem #1: Misalignments in the meaning of implicit cognition. Readers familiar with Daniel Kahneman's best-selling book *Thinking, Fast and Slow* (2011) might

recognize the similarities between Kahneman's distinction between two systems in the mind, System 1 and System 2, and Greenwald and Banaji's (1995) distinction between implicit cognition and explicit cognition. Specifically, Kahneman (2011) defines System 1 as "[operating] automatically and quickly, with little or no effort and no sense of voluntary control" (p. 20), and System 2 as "[allocating] attention to the effortful mental activities that demand it, including complex computations" (p. 21). Indeed, both Kahneman (2011) and Greenwald and Banaji (1995) outline this dual-process typology as a binary distinction between two types of psychological processes that cluster together into two groups, and which underlie human cognitive architecture. These similarities are not surprising, as Banaji (2014) has noted that the study of implicit cognition has largely been the application of Daniel Kahneman and Amos Tversky's work on heuristics and biases to thinking in the social world; and the idea that a dual-process typology underlies human cognitive architecture has only grown more popular since its emergence in the 1970s (Pennycook, Neys, Evans, Stanovich, & Thompson, 2018). Indeed, in recent years this dual-process typology has been canonized in the minds of experts and laypeople alike, and shaped the foundation and direction for empirical and theoretical work on a wide range of topics (for an overview, see Melnikoff & Bargh, 2018).

However, despite the immense popularity of the dual-process typology, it currently lacks any empirically justifiable backing. As reviewed by Melnikoff and Bargh (2018), there is no empirical evidence that processing features (i.e., unawareness-awareness, unconsciousness-consciousness, etc.) exclusively cluster together into two distinct groups (i.e., System 1 and System 2; implicit processes and explicit processes), and there is substantial evidence accumulated over the past 30 years which explicitly falsifies the dual-process typology. This is particularly salient when considering examples where processing features clearly misalign with the System 1 or System 2 group they have been assigned to: for example, humans engage in many skills which are unconscious (System 1) but intentional (System 2), such as driving, typing, or playing a musical instrument (for

more examples, see Melnikoff & Bargh, 2018). To be clear, this is not to say that there is no qualitative distinction between unconscious and conscious processes; rather, it is saying that there is a great deal of evidence that what distinguishes an unconscious process from a conscious process is simply awareness and reportability of the process, not individual processing features such as efficiency, intentionality, controllability, etc. (for reviews, see Baumeister & Bargh, 2014; Melnikoff & Bargh, 2018).

To their credit, Greenwald and Banaji (1995) recognized that awareness and reportability of implicit and explicit processes were of more importance than clusters of individual processing features. In their template for defining specific categories of implicit cognition Greenwald and Banaji (1995) wrote that: “An implicit C is the introspectively unidentified (or inaccurately identified) trace of past experience that mediates R. In this template, C is the label for a construct (such as attitude), and R names the category of responses (such as object-evaluative judgments) assumed to be influenced by that construct.” However, despite this operationalization, the choice to also use the implicit-explicit distinction to capture a wide range of psychological processes clustered together into two distinct groups leaves it with the same fundamental problem as any dual-process typology. Indeed, one might say that this fundamental problem is an implicit feature of Greenwald and Banaji’s (1995) theory of implicit cognition, as evidenced by the continued memetic use of implicit cognition theory—and the measurement tools this theory has inspired—inside and outside of academia (see Greenwald & Banaji, 2017; Mitchell & Tetlock, 2017).

Driving this problem further is the misalignment between Greenwald and Banaji’s (1995) definition of implicit cognition and the outcomes obtained by so-called implicit measures. For example, a series of studies by Hahn, Judd, Hirsh, and Blair (2014) indicate that people actually are aware of the construct(s) responsible for their performance on measures of implicit cognition, as they are able to accurately predict their performance on implicit measures under a variety of testing conditions; even when performance on the

implicit measure diverges from responses to explicit measures. Similar findings using a different implicit measure have also been obtained (see Cummins, Hussey, & Hughes, 2019). Indeed, many theories of implicit cognition fail to account for the findings of research using implicit measures (see Brownstein et al., 2019; Corneille & Stahl, 2019; Gawronski et al., 2006). Despite this misalignment, it is the empirically unsupported and unjustifiable memes about implicit cognition (see Brownstein et al., 2019; Corneille & Stahl, 2019;; Greenwald & Banaji, 2017; Mitchell & Tetlock, 2017) which continue to dominate conversation, research, and policy about implicit biases/attitudes/associations/prejudices/stereotypes (e.g., Comey, 2015; Gladwell, 2007; Google, 2013; Starbucks, 2018).

Problem #2: Imprecise (or inaccurate) language enables poor ideas to propagate. Building upon Problem #1, the decision to use the implicit-explicit terminology, and to embed this into a dual-process typology, has led to a number of memetic ideas surrounding the concept of implicit cognition. Most important to the present discussion is the empirically unjustifiable idea that implicit cognition = unconscious cognition (or System 1) (Greenwald & Banaji, 2017). As noted above, the origins of this idea are found in the paper outlining Greenwald and Banaji's (1995) original theory of implicit cognition, which heavily emphasized the link between implicit cognition and unconscious processes, and the importance of developing new psychological measures capable of measuring implicit cognition. Three years later, Greenwald, McGhee, and Schwartz (1998) introduced the Implicit Association Test (IAT) to the world, claiming it to be one of the first measures capable of measuring individual differences in implicit cognition. Indeed, Greenwald et al. (1998) presented the IAT as if it had decades of empirical support behind it—for example, regarding the outcome of a Race IAT (experiment 3 in their paper) they wrote “At the same time that these findings are encouraging in regard to usefulness of the IAT to measure implicit attitudes, they are discouraging in indicating the pervasiveness of unconscious forms of prejudice” (p. 1475).

Thus, rather than treating the IAT as a brand new measure which had yet to face the scrutiny of peer-review or have its psychometric properties validated, Greenwald et al. (1998) simply made claim of the IAT's measurement capabilities without supporting evidence. Unfortunately, these unsubstantiated claims were rarely challenged in the decade after the IAT's creation (De Houwer, 2006); likely due to these claims aligning with the ideological sympathies and research agendas of researchers interested in the IAT (see Mitchell & Tetlock, 2017).¹ However, the spread of these ideas could have been easily avoided through the use of more honest, scientifically conservative, language.

Unsubstantiated claims such as the above are a pervasive theme underlying the memetic ideas surrounding the concept of implicit cognition. As documented by Mitchell and Tetlock (2017), public dissemination of information about implicit cognition and the IAT began in mass shortly after Greenwald and colleague's (1998) publication, with extensive multimedia campaigns boldly declaring the pervasiveness of implicit biases/attitudes/associations/prejudices/stereotypes (all these labels are necessary due to the imprecise language about implicit cognition), the ability of the IAT to measure these various forms of implicit cognition, and of the link between these various forms of implicit cognition as measured by IAT and various forms of negative social behaviour (e.g., prejudice, discrimination, etc.). Indeed, just as the dual-process typology the concept of implicit cognition and the IAT are founded on has been canonized because of popularity rather than evidence (Melnikoff & Bargh, 2018), so has the validity and reliability of implicit cognition, the IAT, and their relation to human thought and behaviour been canonized in much the same manner (for a review, see Mitchell & Tetlock, 2017).²

¹ Although the topic is not of direct concern in this article, readers may be interested in recent commentary regarding the lack of political diversity in social psychology, and its negative impact on the research questions, methods, and conclusions made by researchers; which ultimately weakens the reliability and validity of social psychological science (see Crawford et al., 2015; Duarte et al., 2015)

² Readers may also be interested in a critical review by Newell and Shanks (2014) examining the role of

This undeserved canonization is further emphasized when considering the theoretical shortcomings of implicit cognition mentioned above, as well as a number of well-supported empirical findings about the IAT, including: (1) Low test-retest reliability of performance on the IAT at the individual level (Cooley & Payne, 2017; Cunningham, Preacher, & Banaji, 2001; Devine, Forscher, Austin, & Cox, 2012; Gawronski, Morrison, Phills, & Galdi, 2017)—as noted by Payne et al. (2017), “the temporal stability of [performance on the IAT] is so low that the same person tested 1 month apart is unlikely to show similar levels of [performance]” (p. 233). (2) Small (or very small) effect sizes of performance on the IAT to predict discriminatory behaviour (Greenwald, Banaji, & Nosek, 2015; Greenwald, Poehlman, Uhlmann, & Banaji, 2009; Oswald et al., 2013, 2015). Additionally, a number of problems implicit to the IAT’s design also demonstrate its shortcomings, specifically: (1) Interpretations of IAT difference scores are based on implicit assumptions about the test-taker’s state of mind and lack an underlying testable model; indeed, it is merely assumed that performance on the IAT reflects a bias/attitude/association/prejudice/stereotype, rather than some alternative cognitive or behavioural process (Fiedler et al., 2006). (2) Thresholds for different levels of bias as measured by the IAT based on the magnitude of IAT difference scores are completely arbitrary (Blanton & Jaccard, 2008). (3) Measures capable of directly testing whether performance on the IAT has anything to do with System 1 processes do not exist (Bar-Anan & Vianello, 2018). (4) Finally, and most importantly, the construct validity of the IAT—and thus, any claims about implicit cognition’s role human thought and behaviour—remain unresolved (De Houwer et al., 2009a, 2009b; cf. Gawronski, LeBel, Peters, & Banse, 2009; cf. Nosek & Greenwald, 2009).

unconscious influences (i.e., awareness) on behaviour. In their review Newell and Shanks (2014) provide a list of necessary criterion for adequate assessments of awareness and its influence on behaviour, note that many studies fail to meet these criterion, and explain how the results of these studies can be explained more parsimoniously without invoking claims of unconscious influence.

Remedies to the Aforementioned Problems

There is no doubt that the psychological phenomena measured by the IAT are both empirically and theoretically interesting, and potentially socially important; however, just as a number of striking findings during the 20th century led to important reevaluations of the nature and measurement of attitudes (for a review, see Greenwald et al., 2002), a number of striking findings during the 21st century have made it necessary to reevaluate the nature and measurement of what is currently termed an implicit bias/attitude/association/prejudice/stereotype. Indeed, Several prominent attitude researchers have expressed similar sentiments:

“Implicit evaluation can be defined as the automatic effect of stimuli on evaluative responses. A major advantage of this definition is that it is neutral with regard to the mental processes and representations that mediate implicit evaluation.” (De Houwer, 2014, p. 342)

“Many publications continue to use implicit and explicit with conceptual meaning (as unconscious vs. conscious), rather than with their better justified empirical (indirect vs. direct) meaning. Even though the present authors find themselves occasionally lapsing to use implicit and explicit as if they had conceptual meaning, they strongly endorse the empirical understanding of the implicit-explicit distinction.” (Greenwald & Banaji, 2017, p. 862)

“Given the potential payoff from better understanding the causes of inequality, and given the real expenditures being made on IAT research and in response to fears caused by the implicit prejudice meme, the implicit prejudice research domain is fertile ground for experimentation with extraordinary science. We have discussed at length what one form of extraordinary science might look like in this domain (Tetlock & Mitchell, 2009), but that is only one possibility.

What is essential is that ground rules for judging success and failure be set ex ante rather than allow contestants to engage in post hoc assimilation of any pattern of results to their preferred theories.” (Mitchell & Tetlock, 2017, p. 187)

“It is time that we as a field come to terms with these issues. With institutions like the World Bank and Institute of Medicine now endorsing our highly speculative and frequently misleading typology, we cannot afford to wait. Luckily, the solution is straightforward: researchers should rigorously explore each feature of a given process one by one, without making assumptions or drawing conclusions about other features that are not being studied. For too long the dual-process typology has obscured the rich diversity of the human mind. Let us embrace that diversity instead.” (Melnikoff & Bargh, 2018, p. 11)

My own suggestion would be to go one step further—removing the term implicit bias/attitude/association/prejudice/stereotype from our vocabulary altogether (outside of theory), and treating the IAT (and other similar implicit measures) as a phenomenon unto itself. That is, the difference score that is calculated to interpret performance on the IAT should simply be referred to as performance on the IAT (e.g., “The manipulation caused significant changes in performance on the IAT.”), rather than, as is common in the literature, as an indicator of an implicit bias/attitude/association/prejudice/stereotype (e.g., “The manipulation caused significant changes in implicit biases/attitudes/associations/prejudices/stereotypes.”) (for a similar recommendation and discussion, see Brownstein et al., 2019).

This recommendation should not be considered controversial; adopting more accurate terminology can only improve the quality and integrity of research, and avoid the spread of empirically unjustifiable ideas throughout the research literature and into the minds of the public, business, and government. Indeed, shifting the current paradigm by treating performance on the IAT as a phenomenon unto itself will also open the door to new,

exciting, research that can help improve our understanding of the cognitive mechanisms that underlie performance on the IAT; particularly so because researchers have absconded from treating the IAT as a dependent measure in experimental research in favour of focusing on research which treats the IAT as a diagnostic end in itself (Fiedler et al., 2006), thus, there are currently many unexplored avenues of inquiry available to researchers interested in the IAT.

Indeed, although social cognition theorists previously hypothesized that implicit biases/attitudes/associations/prejudices/stereotypes as presumed to be measured by the IAT or other implicit measures were stable psychological constructs resilient to change and external influence (e.g., Wilson, Lindsey, & Schooler, 2000); however, a recent meta-analysis by (Forscher et al., 2016) on performance on the IAT found that performance is malleable to a variety of experimental procedures (e.g., associating sets of concepts, invoking goals/motivations, taxing mental resources, etc.), and that most procedures create larger changes on implicit measures than on explicit or behavioural measures (or, indeed, only influence changes on implicit measures). These findings corroborate earlier evidence that performance on the IAT can be changed through internal and external influences (for reviews, see Blair, 2002; Dasgupta, 2009; Gawronski & Bodenhausen, 2006, 2011; Gawronski & Sritharan, 2010; Lai, Hoffman, & Nosek, 2013; Lenton, Bruder, & Sedikides, 2009; Sritharan & Gawronski, 2010), opening up the possibility of experimentally examining the cognitive mechanisms that influence performance on implicit measures, thus creating a better understanding of the cognitive mechanisms underlying or contributing to performance on the IAT.³

³ It may also be worth reexamining past implicit cognition research within the attitude-agnostic framework of treating the IAT as a phenomenon unto itself. Doing so may help generate new hypotheses about the cognitive mechanisms related to performance on the IAT, or uncover new insights which may have been ignored in research that treated the IAT as an attitude measure. Conducting literature reviews, systematic reviews, or meta-analyses under this framework would also be a useful method to keep past research relevant, and perhaps to attract researchers from other areas (e.g., cognitive or experimental psychology)

The purpose of the present research was to test an unexplored area of IAT research—indirectly caused IAT performance change—in order to better understand the cognitive mechanisms underlying performance on the IAT. Previous research in this area has almost exclusively focused on direct forms of influence (e.g., showing images of admired Blacks and disliked Whites prior to a Race-IAT) (for a review, see Gawronski & Bodenhausen, 2006), limiting the scope of the types of cognitive mechanisms which can be attributed to performance on the IAT; namely, simple associational structures (see Karpinski & Hilton, 2001). Extending the possible interpretations of the cognitive mechanisms underlying IAT performance, the present research explores the limits of this topic by testing whether performance on the IAT can be influenced by complex associational structures.

Method

Design

The present research tested the effect priming and conditioning complex associational information had on IAT performance using a simple pre-post experimental design, with the priming manipulation as an independent variable across three conditions. Stimuli on the IAT measures consisted of good and bad words, and the colours red and blue. Stimuli on the priming task consisted of colour images depicting nationalistic displays in Nazi Germany or Canada (experimental conditions), or of people gathered in public places (control condition). The logic of the experiment was that the predominance of the colour red in the Nazi or Canadian images (e.g., on flags or uniforms) would act as an indirect context cue that Red = Bad (Nazi Condition) or Red = Good (Canadian Condition)—because the colour red is an integral feature of both groups' identities—thus influencing performance on the post-test IAT by changing how participants would respond

to contribute to the IAT literature.

to pairings of red with bad or good words. The reliability and validity of the IAT and prime measures were assessed using a set of explicit measures.

Participants

120 students (27 Males, 93 Females; Age, Mean = 20.79; SD = 4.59) students from the University of the Fraser Valley (a Canadian university) participated in exchange for course credit and to enter a raffle for \$100. Data for five of these participants was excluded due to meeting predetermined exclusion criteria. All participants were randomly assigned to a condition using a double-blind procedure, and stimulus pairing orders on the IATs were counterbalanced.

Apparatus

Hardware. Participants completed the experiment on PCs with colour monitors in a computer lab. Each computer terminal was separated by a partition so participants could not see each other's monitors, and each computer was operating on the same hardware and software. Up to 24 computer terminals were available for each session of data-collection. Participants viewed the computer monitor from roughly an arm's length away and responded to experimental prompts using the keyboard or mouse.

Software. The entire experimental procedure, with the exception of the initial welcoming of participants and final debriefing, was programmed and delivered to participants using E-Prime's (Psychology Software Tools, Inc., 2016) E-Studio and E-Run tools, respectively. E-Run sessions for the experiment were loaded prior to participants' arrival. The subject number for each participant's E-Run session was used to randomly assign a participant to a condition, ensuring that assignment to a condition was double-blind.

Random assignment was achieved using the following procedure: (1) Participants

were assigned an ordinal subject number based on order of arrival (i.e., the first participant to arrive was “participant one,” and the last participant to arrive was “participant 120”), which was used to assign participants to a computer terminal (i.e., the first participant to arrive sat at Terminal 1, the second participant to arrive sat at Terminal 2, etc.). (2) Each ordinal subject number was paired with a number randomly generated from a list ranging from 1-120, which was used for each participant’s E-Run subject number (e.g., participant one’s E-Run subject number was 75, participant two’s E-Run subject number was 15, etc.); each value from the randomly generated number list was only used once. This procedure was done prior to data collection.

Materials

Red-blue Implicit Association Test. The pre- and post-test IATs were programmed and presented to participants using the methodology outlined by Greenwald, Nosek, and Banaji (2003), with two minor modifications: First, the test was labelled as a “Decision-making Speed and Accuracy Test,” rather than the “Implicit Association Test.” Second, incorrect responses were indicated with a black “X” instead of the typical red “X”, to avoid confusion with the red-coloured shape items on the IAT. The attribute items on the IAT consisted of a subset selected from those used by Greenwald et al. (1998). Namely, five good-meaning words (*freedom, peace, honest, paradise, love*), and five bad-meaning words (*murder, death, evil, hatred, filth*) were selected. The concept items on the IAT consisted of five blue-coloured shapes, and five red-coloured shapes (*circle, triangle, square, pentagon, star*). The shade and tint of the colours were the same for each item. The RGB hex codes were #0000FF for the blue colour, and #FF0000 for the red colour. A variety of shapes were selected to avoid possible confounds a single colour-shape combination might have, as there is some evidence that certain colour-shape associations exist, and that these associations can influence speed and accuracy results on the IAT (Chen, Tanaka, & Watanabe, 2015; cf. Makin & Wuerger, 2013).

In all other respects the pre- and post-test IATs followed the methodology outlined by Greenwald et al. (2003). Specifically, the IATs consisted of seven different blocks wherein participants had to discriminate the appropriate categorization for attribute (i.e., good or bad) and/or concept (i.e., red or blue) target stimuli presented in the centre of the screen using the “E” and “I” keys on the keyboard. Each block was preceded by a set of instructions concerning the dimensions of the discrimination task, the appropriate key response for categorization, and a notice on how to respond if an error is made. Once participants read the instructions for a block they were instructed by a line of text at the bottom of the screen to press the spacebar proceed with the task in each block. Category reminder labels corresponding to the response keys were also positioned appropriately on the upper left (“E”) and upper right (“I”) sides of the screen during the instructions and trials for each block, as well as instructions at the bottom of the screen on how to respond to errors.

Participants completed the seven blocks of the IAT in the following order: Block 1: Initial target-concept discrimination (blue vs. red). Block 2: Initial evaluative attribute discrimination (good vs. bad). Block 3 and 4: Initial combined task (blue + good vs. red + bad). Block 5: Reversed concept-target discrimination (red vs. blue). Block 6 and 7: Reversed combined task: (red + good vs. blue + bad). For half of the participants in each group the order of the concept-target discrimination was switched such that red vs. blue appeared in Block 2, and blue vs. red appeared in Block 5.

In blocks 1, 2, 3, and 6, participants completed 20 trials. In blocks 4, 5, and 7, participants completed 40 trials. During each trial a single target stimulus was presented in the centre of the screen. The stimulus remained on the screen until the participants made a response. If a correct response was made then the program advanced to the next trial, which involved the stimulus disappearing from the screen without any feedback, and a new stimulus from the next trial appearing after 250 milliseconds. If an incorrect response was made then a black “X” appeared on the screen, at which point participants had to input

the correct response key in order to advance to the next trial. The speed (in milliseconds) and accuracy (correct or incorrect), as well as the stimulus and stimulus categories were recorded for each trial.

Associative priming task. The stimulus items for the priming task consisted of a set of 38 colour images depicting nationalistic displays (mainly neutrally valenced things such as parades, public gatherings, or pictures of military or police) in Nazi Germany or Canada with a predominance of the colour red in the images (experimental conditions), or of people gathered in public places with either no predominant colours or non-red or non-blue predominant colours (control condition) (for examples, see Appendix A). Each trial for the priming task consisted of the following procedure: A randomly selected (without replacement) image was displayed for two-seconds, followed by a randomly selected good word and bad word which participants were asked to select based on which word best reflected their attitude towards the image (taken from the same word-list as the attribute items for the IAT). Good word choices were assigned a value of 50, and bad word choices were assigned a value of -50. The words were presented in the middle-left and middle-right of the computer screen, and the side of the screen the good and bad words were presented on was randomized for each trial in order to prevent participants from engaging in a response set. Participants were instructed to take as much time as they needed to respond to each trial, and their responses were logged in order to be used as an indirect measure of their attitude towards the images. The advantage of this task was that, because the selected images were quite neutral regarding their valence, it could be assumed that participants' responses to the stimuli would be biased based on their prior attitudes and knowledge about Nazis or Canadians, rather than on the content of the pictures themselves (i.e., any response differences between the conditions were likely caused by participant's explicit biases).

Explicit Measures. The explicit measures were presented to participants in the same order as they are presented below. Participants responded to the explicit measures

using either the mouse or keyboard. Participants were also free to skip any questions they did not wish to answer.

Suspicion probe.

The suspicion probe consisted of a small battery of open- and close-ended questions intended to assess participant's thoughts about the experiment, whether they felt they were being deceived and whether this influenced their behaviour, whether they felt certain reactions were expected of them, and whether they had any information about the study prior to participating other than that provided by the experimenter. The purpose of these questions was to get a rough idea of the quality of data obtained from the pre- and post IATs and the priming task, and of participant's thoughts and feelings towards those tasks.

Colour preference likert scale.

The colour-preference likert scale consisted of a 7-point scale asking participants the strength of their preference of Red relative to Blue (or vice versa) with the prompt "Which statement best describes you?" The degrees of strength were expressed as strongly, moderately, and slightly (i.e., *I strongly/moderately/slightly prefer Red to Blue*; or, *I strongly/moderately/slightly prefer Blue to Red*) with a neutral middle point (i.e., *I like Red and Blue equally*). Responses were scored on a scale from -99 (*I strongly prefer Blue to Red*) to 99 (*I strongly prefer Red to Blue*), with scale intervals of 33 and 0 as a neutral middle point (*I like Red and Blue equally*).

Feeling thermometers.

On the feeling thermometers, participants were asked to describe their general level of warmth or coolness toward the colour blue, the colour red, and toward Nazis/Canadians/People (condition dependent) by moving a continuous slider to the appropriate position on an illustration of a thermometer. The thermometer was labelled with tick-marks at 5-degree intervals and anchors with the words extremely cold at the lowest value, neither warm nor cold at the neutral point, and extremely warm at the

highest value. Scores ranged from -50 (*extremely cold*) to 50 (*extremely warm*), with 0 as a neutral middle point (*neither warm nor cold*).

Red-blue Attitude change question.

Because this experiment did not include pre-test measures of participants' explicit attitudes towards the colours red and blue, two questions asking participants "Did viewing the pictures of Nazis/Canadians/People (condition dependent) change your attitude towards the colour Red/Blue (question 1/question 2)?" were included as a very crude measure of possible explicit attitude change. Participants were able to indicate whether their attitude changed (yes or no) and the direction of the change (positive or negative). This question was chosen mainly for pragmatic reasons, as it was unknown whether including the colour preference likert scale and red and blue feeling thermometers as pre-test measures would interfere with the IATs or priming task in some unforeseen way, and the present author did not have the available resources or participant pool to first conduct a version of the present experiment with counterbalanced pre- or post-test explicit colour measures, and then conduct a version with pre- and post-test explicit colour measures.

Identity-condition congruency question.

This question simply asked participants, yes or no, whether they identified as a Nazi/Canadian/Person (condition dependent); both because self-identity seems to influence IAT performance when one's self-identity is related to content on the IAT (e.g., Greenwald et al., 1998), and because this served as a measure of whether participants considered the groups depicted in the priming task as in-group or out-group members.

Demographic questions.

Demographic questions consisted of a question asking participants if they were colour blind for red and/or blue, and basic demographic questions (age, gender identity). The colour blindness question was included as it was assumed that colour blindness for red

might exclude a participant from being effected by the experimental manipulation.

Procedure

Upon arrival the experimenter greeted each participant and asked them to sit at their assigned computer terminal. Each computer terminal was set to the title page of the experiment's electronic informed-consent form, displaying a dummy title for the experiment ("The Effects of Image Valence on Response Speed and Accuracy"), basic institutional information, and the date and time of the data-collection session. Once all participants were seated (or once the scheduled time for the data-collection session had arrived) the experimenter read a short introduction to participants, introducing themselves, welcoming participants to the experiment, outlining behavioural expectations for the experiment (e.g., "Do not speak to other participants or try to look at each other's screens." "Please turn off and put away your cell-phones."), and then instructing participants to press the space bar on their keyboard, allowing them to read the electronic informed-consent form. Please note: From this point onwards the experimenter had no further interaction with participants until all participants had completed the experiment. All remaining information and instructions about the experiment was given to participants via text and images on the computer screen.

After reading the informed-consent participants were instructed to either notify the researcher if they wished to withdraw, or to press the enter key on their keyboard if they wished to participate. Participants who chose to withdraw were still granted course credit, but were not entered into the raffle. Participants who chose to participate completed the experiment's tasks in the following order (for the design of each task please read the materials section): (1) Pre-test Red-Blue IAT, (2) Free-association priming task, (3) Post-test Red-Blue IAT, (4) Explicit measures. Each task was separated by a prompt indicating to participants that they had completed the current task, and to press the spacebar to continue to the next task. Additionally, each task was introduced to

participants with a set of instructions outlining the content of the task, and of how to respond to the task (for exact wordings, see Appendix B).

Participants were introduced to both IATs as a decision-making speed and accuracy test in which they were to sort items from four different categories as quickly as they could; all stimuli and their respective categories were introduced to participants prior to beginning the task. The instructions for each block of the IAT simply told participants which response keys to use for each categorization, and reminded participants to respond as quickly and accurately as they could. On Block 5, where the categories switched sides, the switch was explicitly noted in the instructions. For the priming task participants were told that they would see pictures of (condition dependent), “Nazis celebrating their rise to power,” “Canadians celebrating their culture,” or “People at public events,” and to reflect on their attitudes towards each picture and select the word that best reflected their attitude for each trial. Participants responded using the mouse, and were told to take as much time as they needed.

Once each participant completed the experiment the program asked them to notify the experimenter; after all participants completed the experiment they were verbally debriefed and thanked for their participation by the experimenter, and their personal information was collected in order to grant them course credits and enter them into the raffle.

Results

Data analysis was conducted in RStudio (RStudio, Inc., 2015) to ensure all analyses were reproducible. An Rmarkdown version of this manuscript containing the analysis scripts can be accessed at (link to be added later).

Analysis of Explicit Measures

Colour Preference. Feeling thermometer difference scores were calculated by subtracting each participant's response to the blue feeling thermometer from their response to the red feeling thermometer, creating an overall colour preference feeling thermometer variable, where positive values indicated a preference for red relative to blue, and negative values indicated a preference for blue relative to red. On a scale from absolute blue preference (-50) to absolute red preference (50), none of the conditions showed a strong colour preference, Nazi (Mean = 5.88; SD = 18.71), Canadian (Mean = 1.84; SD = 17.51), Control (Mean = 0.60; SD = 24.28), and there were no significant differences in colour preference between the conditions, $F(2, 112) = 0.70$, $MSE = 416.82$, $p = .499$, $\hat{\eta}_G^2 = .012$. Similar results were obtained for the colour preference likert, Nazi (Mean = -4.38; SD = 27.58), Canadian (Mean = NA; SD = NA), Control (Mean = NA; SD = NA), $F(2, 107) = 0.27$, $MSE = 952.46$, $p = .762$, $\hat{\eta}_G^2 = .005$.

A correlational analysis was then conducted between the colour preference feeling thermometer and participant's responses to the likert-scale colour preference question, revealing a significant positive correlation, $r = .40$, 95% CI [.23, .54], $t(108) = 4.51$, $p < .001$, indicating that the two measures tapped into a similar construct. Because of this, the two measures were combined into an overall colour preference variable.

Attitudes Towards the Prime. Feeling thermometer scores towards Nazis/Canadians/People trended in the expected direction. On a scale from extremely cold (-50) to extremely warm (50), participants felt extremely cold towards Nazis (Mean = -39.08; SD = 21.38), moderately warm towards Canadians (Mean = 28.79; SD = 17.48), and neutral towards People (Mean = 0.69; SD = 27.40), indicating that participants held negative attitudes towards Nazis, positive attitudes towards Canadians, and neutral attitudes towards People (likely because the category of people is extremely broad and vague, prompting an equally vague response; whereas the categories of Nazi or Canadian

are salient and familiar, prompting the expected responses). The difference between these attitudes was significant, $F(2, 112) = 87.21$, $MSE = 506.78$, $p < .001$, $\hat{\eta}_G^2 = .609$.

Responses on the free-association priming task also trended in the expected direction. On a scale from only bad-word choices (-50) to only good-word choices (50), participants felt that bad words best represented Nazis (Mean = -33.86; SD = 36.80), that good words best represented Canadians (Mean = 37.33; SD = 33.28), and good words best represented People, (Mean = 29.49; SD = 40.39), indicating that participants responded to the images using their a priori biases and knowledge about Nazis or Canadians. The choice of good words best representing People likely indicated that good choices (freedom, peace, honest, paradise, love) were a more appropriate default response than bad choices (murder, death, evil, hatred, filth). The differences between these responses were significant, $F(2, 112) = 207.32$, $MSE = 280.13$, $p < .001$, $\hat{\eta}_G^2 = .787$.

A correlational analysis was then conducted between the prime feeling thermometer and free-association task, revealing a significant positive correlation, $r = .77$, 95% CI [.68, .83], $t(113) = 12.76$, $p < .001$, indicating that the two measures tapped into a similar construct. Because of this the two measures were combined into an overall prime attitude variable.

Analysis of the IATs

Analysis of Exclusion Criteria. In accordance with Greenwald, Nosek, and Banaji's (2003) improved scoring algorithm, data from four participants was eliminated due to having response latencies of less than 300 milliseconds on more than 10% of the trials.

Performance on the IATs. IAT performance was scored using the IAT scoring algorithm recommended by Greenwald et al. (2003). In the context of the present experiment, positive IAT scores indicated faster and more accurate responses when the categories Red and Good were paired together than when the categories Red and Bad were

Figure 1. Pre- and post-test comparisons of IAT performance between the conditions.

paired together, and negative IAT scores indicated faster and more accurate responses when the categories Red and Bad were paired together than when the categories Red and Good were paired together. Performance on the pre-test IAT revealed some mean differences between the groups, Nazi (Mean = -0.14; SD = 0.53), Canadian (Mean = 0.04; SD = 0.46), Control (Mean = -0.15; SD = 0.48), however, these differences were not significant, $F(2, 112) = 1.85$, $MSE = 0.24$, $p = .163$, $\hat{\eta}_G^2 = .032$.

Did the Prime Affect IAT Performance? A 3 (condition: Nazi, Canadian, Control) X 2 (time: pre-test vs. post-test) ANOVA with repeated measures on time was conducted on IAT Effect scores in order to assess the effect of the priming task on IAT performance (Figure 1).. This analysis revealed no significant differences between Condition, $F(2, 112) = 0.83$, $MSE = 0.26$, $p = .441$, $\hat{\eta}_G^2 = .010$; Time, $F(1, 112) = 3.95$,

$MSE = 0.11$, $p = .049$, $\hat{\eta}_G^2 = .010$; or the Condition X Time interaction, $F(2, 112) = 2.20$, $MSE = 0.11$, $p = .116$, $\hat{\eta}_G^2 = .011$, indicating that the priming task did not successfully influence participant's performance on the post-test Red-Blue IAT. Indeed, performance on the post-test Red-Blue IAT was near-identical between conditions, with mean IAT scores at zero for all conditions, Nazi (Mean = 0.02; SD = 0.32), Canadian (Mean = 0.00; SD = 0.34), Control (Mean = -0.01; SD = 0.42), indicating that performance was not biased towards either possible stimulus pairing (Red + Good or Red + Bad) on the post-test IAT. However, it should be noted that responses on the priming task were biased in the expected direction, Nazis (Mean = -33.86; SD = 36.80), Canadians (Mean = 37.33; SD = 33.28), People, (Mean = 29.49; SD = 40.39), indicating that the prime's lack of influence on participants' performance on the post-test Red-Blue IAT was not due to a weak response to the prime.

Subset analysis.

The same 3 X 2 Anova was conducted excluding data from participants who responded positively overall to the Nazi stimuli or negatively overall to the Canadian stimuli; this analysis also revealed no significant differences between Condition, $F(2, 105) = 0.82$, $MSE = 0.25$, $p = .445$, $\hat{\eta}_G^2 = .011$; Time, $F(1, 105) = 4.10$, $MSE = 0.11$, $p = .045$, $\hat{\eta}_G^2 = .012$; or the Condition X Time interaction, $F(2, 105) = 2.28$, $MSE = 0.11$, $p = .107$, $\hat{\eta}_G^2 = .013$.

Analysis of Implicit-Explicit Correlations

Correlational analyses were conducted between implicit and explicit attitude measures in order to assess the reliability and validity of the IAT and prime measures (Table 1). These analyses revealed significant positive correlations between the pre- and post-test IATs, $r(113) = 0.43$, $p < .001$, indicating that performance on both tests was related but not identical; between the pre-test IAT and the colour preference likert, $r(108) = 0.27$, $p = .004$, and the post-test IAT and the colour preference likert, $r(108) = 0.35$, $p < .001$,

Table 1

Matrix of correlations between implicit and explicit measures.

	1	2	3	4	5	6
1. IAT.pre	-					
2. IAT.post	0.43***	-				
3. LIKERT.colour	0.27**	0.35***	-			
4. FT.colour	0.07	0.1	0.4***	-		
5. FT.prime	0.07	-0.16	-0.13	-0.14	-	
6. RESP.prime	0.05	-0.12	-0.1	-0.08	0.77***	-

indicating a possible association between participants' attitudes towards the colours and their performance on the IAT, or simply reflecting the neutral outcomes of both the measures; between the colour preference likert and the colour feeling thermometers, $r(108) = 0.40$, $p < .001$, indicating that the two measures likely captured similar attitudes about the colours red and blue; and between responses to the prime task and responses to the prime feeling thermometer, $r(113) = 0.77$, $p < .001$, indicating that the measures captured participants' attitudes towards the prime images with a high degree of overlap.

There were no significant correlations between attitudes for the prime images and any of the implicit or explicit measures, likely indicating that participants did not make an association between the prime images and the stimuli on the IAT. Indeed, an analysis of the open-ended stimulus probe question "In your own words, what was the present study about?" revealed that only seven participants made an explicit connection between the prime images and the stimuli on the IAT, and of these, only one participant's performance on the IAT moved in the expected direction (i.e., they were in the Nazi condition and their performance became more biased to Red + Bad on the post-test).

Discussion

Although the theoretical assumptions thought necessary for the experimental manipulation to have an effect were met—namely that participants felt negatively towards Nazis or positively towards Canadians, respectively—it appears that participants' explicit attitudes towards the priming images did not influence their performance on the post-test Red-Blue IAT. Although this result is not surprising, it is of theoretical interest in light of other research which has demonstrated that performance on the IAT can be influenced through evaluative conditioning and changes in context cues (for a review, see Gawronski & Bodenhausen, 2006).

The present experiment utilized both these mechanisms in its priming task, but set itself apart from prior experiments in the relationship that existed between IAT and prime items. Where prior experiments utilized a prime that mirrored the IAT content exactly (e.g., showing images of admired Blacks and disliked Whites prior to a Race-IAT), the present experiment utilized a holistically represented prime (Nazi, Canadian) and an IAT that measured a singular component of that prime (the colour red). This may indicate that performance on the IAT may be insensitive to complex relational information (i.e., Nazi = Red, therefore Red = Bad). Such a proposition seems reasonable when considering other research that has demonstrated participants respond to the IAT at the categorical level rather than the stimuli level (De Houwer, 2001), in which case it would be expected that IAT performance would only change under conditions where perceptions of the categories themselves change (e.g., if the Red category was relabelled as Nazi on the post-test IAT).

Alternatively, it is possible that the design of the priming task was too indirect to create an effect, as it was designed to assess whether a performance change would happen spontaneously without any explicit guidance from the experimenter. Although this was the purpose of the experiment, it is feasible to assume that directing participants' attention towards the association between the colour red and the images from the priming task may

have changed the outcome of the experiment; particularly because of participants' strongly valenced responses to the priming images, and because previous experiments have demonstrated that the threshold for a performance change on the IAT to occur is quite undemanding. For example, Karpinski and Hilton (2001) found that simply exposing participants to repeated pairings of the word elderly with positive words on a Young-Old IAT was sufficient to significantly reduce a performance bias favouring a Young + Good pairing over an Old + Good pairing. Indeed, this is the typical finding in IAT performance change research (Forscher et al., 2016). Thus, it is possible that IAT performance change might only occur in situations where the mechanism of change is sufficiently salient, indicating that the psychological constructs measured by the IAT are possibly centred around simple learned-association structures and/or basic sensori-motor processes. Alternatively, the saliency of environmental cues may be used by participants in a more elaborate way as a cognitive anchor (i.e., a test-taking strategy) allowing them to perform more quickly and accurately on certain blocks of the IAT.

Such possibilities might also explain why performance on the IAT is so unstable (e.g., Cooley & Payne, 2017; Cunningham et al., 2001; Devine et al., 2012; Gawronski et al., 2017), as many of the topics IAT researchers are interested in (e.g., race) are highly salient, making these stimuli easier to respond to using data-driven, bottom-up processes, or using performance strategies that take advantage the saliency (i.e., using a cognitive heuristic to respond to stimuli more quickly and accurately). Indeed, because the IAT merely instructs participants to respond as quickly and accurately as they can (as opposed to responding in an unbiased manner), it is possible that performance differentials on different blocks of the IAT reflect heuristics of different strength and saliency being available to participants during certain blocks of the test, creating a biased IAT score. In other words, certain pairings on the IAT may elicit better strategies for responding more quickly and accurately than other pairings, and the saliency of test-related environmental cues may aid in the strength or weakness of some strategies over others (for a related theory, see Payne et al.,

2017). Thus, perhaps encountering positive or negative exemplars of a category is able to shift people's schemas about a category sufficiently enough to influence performance on the IAT (whether by changing attitudes, available test-taking strategies, or some other cognitive mechanism; the exact cognitive mechanisms are open to speculation). However, it should be noted that much of the research studying performance changes on the IAT is confined to measurement(s) from a single experimental session, which, consequently, only provides evidence of short-term change (Forscher et al., 2016). Thus, it cannot currently be determined what types of events (internal or external) create temporary, long-lasting, or permanent, performance changes on the IAT.

Limitations

Because there is little to no research about indirectly caused IAT performance change, it is difficult to make any conclusive statements about the present experiment. Indeed, the present experiment is best thought of as exploratory, and would best be used to guide future research in this area. One shortcoming of the present experiment that might be useful to address in future studies is the implementation of explicit measures. Specifically, future experiments would likely benefit from making a subset of their explicit measures into pre- and post-measures, as opposed to only post-measures. In the context of the present experiment, it would have been useful to see whether explicit attitudes towards the colours red and blue were influenced by the images of Nazis/Canadians/People. Although there was a question asking participants if seeing these images changed their attitude towards the colours red and blue, this was a very crude measure, and it was difficult to gauge whether there was an effect, or whether participants responses were arbitrary; particularly because these responses would have relied on post-hoc reasoning, and the wording of the question may have elicited certain responses from participants.

Conclusion

Although the present experiment did not discover any new cognitive phenomenon that influence performance on the IAT, the results seem to reaffirm older claims that the psychological phenomena measured by the IAT are centred around simple learned-association structures. If the psychological phenomena measured by the IAT were centred around more complex learned-association structures, it would have been expected that the present experiment's priming task would have elicited a response, as it utilized evaluative conditioning and changes in context cues in a manner that was, arguably, quite powerful. Specifically, the colour red was conditioned to co-occur with either Nazis or Canadians, which participants elicited strong evaluative responses to; and participants were presented with the perceptually obvious context cue that the colour red is a fundamental feature of Nazi or Canadian imagery. Considering many other IAT performance change experiments have used much less powerful priming methods with success (e.g., Karpinski & Hilton's [2001] word pairing task discussed earlier), the most parsimonious conclusion regarding the present experiment's results is that the cognitive mechanisms underlying IAT performance are insensitive to complex associations.

References

- Banaji, M. R. (2014). On Kahneman: A Reality Club Discussion on the Work of Daniel Kahneman. *Edge*. Retrieved from:
https://www.edge.org/conversation/daniel_kahneman-on-kahneman.
- Bar-Anan, Y., & Vianello, M. (2018). A multi-method multi-trait test of the dual-attitude perspective. *Journal of Experimental Psychology: General*, 147(8), 1264–1272.
doi:10.1037/xge0000383
- Baumeister, R. F., & Bargh, J. A. (2014). Conscious and unconscious: Toward an integrative understanding of human mental life and action. In *Dual-process theories of the social mind* (pp. 35–49). Guilford Press.
- Blair, I. V. (2002). The malleability of automatic stereotypes and prejudice. *Personality and Social Psychology Review*, 6(3), 242–261. doi:10.1207/S15327957PSPR0603_8
- Blanton, H., & Jaccard, J. (2008). Unconscious racism: A concept in pursuit of a measure. *Annual Review of Sociology*, 34, 277–297. doi:10.1146/annurev.soc.33.040406.131632
- Brownstein, M., Madva, A., & Gawronski, B. (2019). What do implicit measures measure? *Wiley Interdisciplinary Reviews: Cognitive Science*, e1501. doi:10.1002/wcs.1501
- Chen, N., Tanaka, K., & Watanabe, K. (2015). Color-shape associations revealed with implicit association tests. *PloS One*, 10(1), e0116954.
doi:10.1371/journal.pone.0116954
- Comey, J. B. (2015). Hard Truths: Law Enforcement and Race. Retrieved from:
<https://www.fbi.gov/news/speeches/hard-truths-law-enforcement-and-race>.
- Conrey, F. R., Sherman, J. W., Gawronski, B., Hugenberg, K., & Groom, C. J. (2005). Separating multiple processes in implicit social cognition: The quad model of implicit task performance. *Journal of Personality and Social Psychology*, 89(4), 469–487. doi:10.1037/0022-3514.89.4.469

- Cooley, E., & Payne, B. K. (2017). Using Groups to Measure Intergroup Prejudice. *Personality and Social Psychology Bulletin*, 43(1), 46–59.
doi:10.1177/0146167216675331
- Corneille, O., & Stahl, C. (2019). Associative Attitude Learning: A Closer Look at Evidence and How It Relates to Attitude Models. *Personality and Social Psychology Review: An Official Journal of the Society for Personality and Social Psychology, Inc*, 23(2), 161–189. doi:10.1177/1088868318763261
- Costantini, G., Perugini, M., Dentale, F., Barbaranelli, C., Alessandri, G., Vecchione, M., & Caprara, G. V. (2019). Assessing positive orientation with the Implicit Association Test. *European Journal of Psychological Assessment*, 35(1), 109–116.
doi:10.1027/1015-5759/a000362
- Crawford, J. T., Duarte, J. L., Haidt, J., Jussim, L., Stern, C., & Tetlock, P. E. (2015). It may be harder than we thought, but political diversity will (still) improve social psychological science: Authors' Response. *Behavioral and Brain Sciences*, 38.
- Cummins, J., Hussey, I., & Hughes, S. (2019). The AMPeror's New Clothes: Performance on the Affect Misattribution Procedure is Mainly Driven by Awareness of Influence of the Primes.
- Cunningham, W. A., Preacher, K. J., & Banaji, M. R. (2001). Implicit attitude measures: Consistency, stability, and convergent validity. *Psychological Science*, 12(2), 163–170. doi:10.1111/1467-9280.00328
- Dasgupta, N. (2009). Mechanisms underlying the malleability of implicit prejudice and stereotypes: The role of automaticity and cognitive control. In T. D. Nelson (Ed.), *Handbook of prejudice, stereotyping, and discrimination*. (pp. 267–284). New York, NY: Psychology Press.
- De Houwer, J. (2001). A structural and process analysis of the Implicit Association Test. *Journal of Experimental Social Psychology*, 37(6), 443–451.

doi:10.1006/jesp.2000.1464

- De Houwer, J. (2006). What Are Implicit Measures and Why Are We Using Them? In R. W. Wiers & A. W. Stacy (Eds.), *Handbook of implicit cognition and addiction*. (pp. 11–28). Thousand Oaks, CA: Sage Publications, Inc. doi:10.4135/9781412976237.n2
- De Houwer, J. (2014). A propositional model of implicit evaluation. *Social and Personality Psychology Compass*, 8(7), 342–353. doi:10.1111/spc3.12111
- De Houwer, J., Teige-Mocigemba, S., Spruyt, A., & Moors, A. (2009a). Implicit measures: A normative analysis and review. *Psychological Bulletin*, 135(3), 347–368. doi:10.1037/a0014211
- De Houwer, J., Teige-Mocigemba, S., Spruyt, A., & Moors, A. (2009b). Theoretical claims necessitate basic research: Reply to Gawronski, Lebel, Peters, and Banse (2009) and Nosek and Greenwald (2009). *Psychological Bulletin*, 135(3), 377–379. doi:10.1037/a0015328
- Devine, P. G., Forscher, P. S., Austin, A. J., & Cox, W. T. L. (2012). Long-term reduction in implicit race bias: A prejudice habit-breaking intervention. *Journal of Experimental Social Psychology*, 48(6), 1267–1278. doi:10.1016/j.jesp.2012.06.003
- Duarte, J. L., Crawford, J. T., Stern, C., Haidt, J., Jussim, L., & Tetlock, P. E. (2015). Political diversity will improve social psychological science. *The Behavioral and Brain Sciences*, 38, e130. doi:10.1017/S0140525X14000430
- Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and uses. *Annual Review of Psychology*, 54, 297–327. doi:10.1146/annurev.psych.54.101601.145225
- Fiedler, K., Messner, C., & Bluemke, M. (2006). Unresolved problems with the “I”, the “A”, and the “T”: A logical and psychometric critique of the Implicit Association Test (IAT). *European Review of Social Psychology*, 17(1), 74–147.

doi:10.1080/10463280600681248

Forscher, P. S., Lai, C. K., Axt, J., Ebersole, C. R., Herman, M., Devine, P. G., & Nosek, B. A. (2016). A Meta-Analysis of Procedures to Change Implicit Measures.

PsyArXiv. doi:10.31234/osf.io/dv8tu

Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change.

Psychological Bulletin, 132(5), 692–731. doi:10.1037/0033-2909.132.5.692

Gawronski, B., & Bodenhausen, G. V. (2011). The associative-propositional evaluation model: Theory, evidence, and open questions. In J. M. Olson & M. P. Zanna (Eds.), *Advances in experimental social psychology*, Vol 44. (pp. 59–127). San Diego, CA: Academic Press. doi:10.1016/B978-0-12-385522-0.00002-0

Gawronski, B., Hofmann, W., & Wilbur, C. J. (2006). Are "implicit" attitudes unconscious? *Consciousness and Cognition*, 15(3), 485–499. doi:10.1016/j.concog.2005.11.007

Gawronski, B., LeBel, E. P., Peters, K. R., & Banse, R. (2009). Methodological issues in the validation of implicit measures: Comment on De Houwer, Teige-Mocigemba, Spruyt, and Moors (2009). *Psychological Bulletin*, 135(3), 369–372.

doi:10.1037/a0014820

Gawronski, B., Morrison, M., Phills, C. E., & Galdi, S. (2017). Temporal stability of implicit and explicit measures: A longitudinal analysis. *Personality and Social Psychology Bulletin*, 43(3), 300–312. doi:10.1177/0146167216684131

Gawronski, B., & Sritharan, R. (2010). Formation, change, and contextualization of mental associations: Determinants and principles of variations in implicit measures. In B. Gawronski & B. K. Payne (Eds.), *Handbook of implicit social cognition: Measurement, theory, and applications*. (pp. 216–240). New York, NY: The Guilford Press.

- Gladwell, M. (2007). *Blink: The Power of Thinking Without Thinking* (1 edition.). Back Bay Books.
- Google. (2013). Unconscious Bias @ Work (Video Series). Retrieved from: <https://rework.withgoogle.com/guides/unbiasing-raise-awareness/steps/introduction/>.
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, 102(1), 4–27.
- Greenwald, A. G., & Banaji, M. R. (2017). The implicit revolution: Reconceiving the relation between conscious and unconscious. *American Psychologist*, 72(9), 861–871. doi:10.1037/amp0000238
- Greenwald, A. G., Banaji, M. R., & Nosek, B. A. (2015). Statistically small effects of the Implicit Association Test can have societally large effects. *Journal of Personality and Social Psychology*, 108(4), 553–561. doi:10.1037/pspa0000016
- Greenwald, A. G., Banaji, M. R., Rudman, L. A., Farnham, S. D., Nosek, B. A., & Mellott, D. S. (2002). A unified theory of implicit attitudes, stereotypes, self-esteem, and self-concept. *Psychological Review*, 109(1), 3–25. doi:10.1037/0033-295X.109.1.3
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology*, 74(6), 1464–1480.
- Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the implicit association test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology*, 85(2), 197–216.
- Greenwald, A. G., Poehlman, T. A., Uhlmann, E. L., & Banaji, M. R. (2009). Understanding and using the Implicit Association Test: III Meta-analysis of predictive validity. *Journal of Personality and Social Psychology*, 97(1), 17–41.

doi:10.1037/a0015575

Hahn, A., Judd, C. M., Hirsh, H. K., & Blair, I. V. (2014). Awareness of implicit attitudes. *Journal of Experimental Psychology. General*, 143(3), 1369–1392.

doi:10.1037/a0035028

Johnson, T. J., Ellison, A. M., Dalembert, G., Fowler, J., Dhingra, M., Shaw, K., & Ibrahim, S. (2017). Implicit bias in pediatric academic medicine. *Journal of the National Medical Association*, 109(3), 156–163. doi:10.1016/j.jnma.2017.03.003

Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Farrar, Straus and Giroux.

Karpinski, A., & Hilton, J. L. (2001). Attitudes and the Implicit Association Test. *Journal of Personality and Social Psychology*, 81(5), 774–788.

doi:10.1037/0022-3514.81.5.774

Kühnen, U., Schießl, M., Bauer, N., Paulig, N., Pöhlmann, C., & Schmidhals, K. (2001). How robust is the IAT? Measuring and manipulating implicit attitudes of East- and West-Germans. *Zeitschrift Für Experimentelle Psychologie*, 48(2), 135–144.

doi:10.1026//0949-3946.48.2.135

Lai, C. K., Hoffman, K. M., & Nosek, B. A. (2013). Reducing implicit prejudice. *Social and Personality Psychology Compass*, 7(5), 315–330. doi:10.1111/spc3.12023

Lenton, A. P., Bruder, M., & Sedikides, C. (2009). A meta-analysis on the malleability of automatic gender stereotypes. *Psychology of Women Quarterly*, 33(2), 183–196.

doi:10.1111/j.1471-6402.2009.01488.x

Makin, A. D. J., & Wuerger, S. M. (2013). The IAT shows no evidence for Kandinsky's color-shape associations. *Frontiers in Psychology*, 4. doi:10.3389/fpsyg.2013.00616

Malinen, S., & Johnston, L. (2013). Workplace ageism: Discovering hidden bias.

Experimental Aging Research, 39(4), 445–465. doi:10.1080/0361073X.2013.808111

Melnikoff, D. E., & Bargh, J. A. (2018). The Mythical Number Two. *Trends in Cognitive*

- Sciences*, 22(4), 280–293. doi:10.1016/j.tics.2018.02.001
- Mitchell, G., & Tetlock, P. E. (2017). Popularity as a Poor Proxy for Utility: The Case of Implicit Prejudice. In S. O. Lilienfeld & I. D. Waldman (Eds.), *Psychological Science Under Scrutiny* (pp. 164–195). Hoboken, NJ, USA: John Wiley & Sons, Inc. doi:10.1002/9781119095910.ch10
- Newell, B. R., & Shanks, D. R. (2014). Unconscious influences on decision making: A critical review. *Behavioral and Brain Sciences*, 37(1), 1–19. doi:10.1017/S0140525X12003214
- Nosek, B. A., & Greenwald, A. G. (2009). (Part of) the case for a pragmatic approach to validity: Comment on De Houwer, Teige-Mocigemba, Spruyt, and Moors (2009). *Psychological Bulletin*, 135(3), 373–376. doi:10.1037/a0015047
- Oswald, F. L., Mitchell, G., Blanton, H., Jaccard, J., & Tetlock, P. E. (2013). Predicting ethnic and racial discrimination: A meta-analysis of IAT criterion studies. *Journal of Personality and Social Psychology*, 105(2), 171–192. doi:10.1037/a0032734
- Oswald, F. L., Mitchell, G., Blanton, H., Jaccard, J., & Tetlock, P. E. (2015). Using the IAT to predict ethnic and racial discrimination: Small effect sizes of unknown societal significance. *Journal of Personality and Social Psychology*, 108(4), 562–571. doi:10.1037/pspa0000023
- Payne, B. K., Vuletich, H. A., & Lundberg, K. B. (2017). The Bias of Crowds: How Implicit Bias Bridges Personal and Systemic Prejudice. *Psychological Inquiry*, 28(4), 233–248. doi:10.1080/1047840X.2017.1335568
- Pennycook, G., Neys, W. D., Evans, J. S. B., Stanovich, K. E., & Thompson, V. A. (2018). The mythical dual-process typology. *Trends in Cognitive Sciences*, 22(8), 667–668. doi:10.1016/j.tics.2018.04.008
- Psychology Software Tools, Inc. (2016). E-Prime 3.0. Psychology Software Tools, Inc.

RStudio, Inc. (2015). RStudio: Integrated Development for R. Boston, MA: Rstudio, Inc.

Sritharan, R., & Gawronski, B. (2010). Changing implicit and explicit prejudice: Insights from the associative-propositional evaluation model. *Social Psychology, 41*(3), 113–123. doi:10.1027/1864-9335/a000017

Starbucks. (2018). The Third Place: Our Commitment, Renewed. Retrieved from: <https://stories.starbucks.com/stories/2018/thethirdplace/>.

Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review, 107*(1), 101–126.

Appendix A

Canadian Images:

Nazi Images:

Control Images:

Appendix B

Note: If replicating this experiment, I would recommend making all of these questions open-ended responses, as the multiple-choice responses did not provide any information that was particularly useful for guiding analysis.

Question 1: In your own words, what was the present study about? -

[OPEN-ENDED RESPONSE]

Question 2: Did you believe, at any time, that the study dealt with anything other than what the experimenter described to you? - Yes - No

Question 3: If you believed that the study dealt with anything other than what the experimenter described, did this affect your behaviour? - Yes - No - Not Applicable

Question 4: Did you feel that certain reactions were expected from you at any time? - Yes - No

Question 5: Did you have any information about the study, other than that provided by the experimenter, prior to participating? - Yes - No