THREE ESSAYS ON THE RELATIONSHIP BETWEEN IMPLICIT ATTITUDES AND POLITICAL BEHAVIOR

BY VINCENT GRECO

A dissertation submitted to the

Graduate School—New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Political Science

Written under the direction of

Richard Lau

and approved by

_____

_____

_____

_____

New Brunswick, New Jersey

January, 2015

ABSTRACT OF THE DISSERTATION

Three Essays on the Relationship between Implicit Attitudes and Political Behavior

by Vincent Greco

Dissertation Director: Richard Lau

Considerable evidence supports the view that information processing often operates in an automatic, less conscious mode of cognition (Lodge & Taber, 2013). Such automatic or "implicit attitudes" have been shown to be a strong predictor of a variety of political outcomes (Perez, 2013). Despite such evidence, it remains unclear the mechanism linking implicit attitudes to political decision-making. This dissertation project attempts to shed light on this issue by examining the impact of implicit attitudes on political information search and candidate evaluation. In my experimental design, I use both the implicit association test (IAT) and the dynamic information board methodology (Lau & Redlawsk, 2006) to directly assess the effect of implicit racial attitudes on how subjects search for information and evaluate political candidates in the context of a mock presidential campaign. I also address a key criticism in this literature by exploring under what conditions implicit attitudes are likely to have stronger and weaker effects on political behavior. To do this, I include a set of manipulations that vary a key feature of the information environment (high vs. low cognitive load) and subjects' emotional state (positive vs. negative). The results from this project provide a direct assessment of both implicit attitudes' impact on information search and candidate evaluation as well as potential moderating factors. Turning to the results, my main finding is that implicit

racial bias is a strong predictor of subjects' information search patterns, which in turn impact candidate evaluation and vote choice. However, as I demonstrate in subsequent chapters, the relationship between implicit bias and voter decision-making is moderated by changes to the information environment. The results from this project provide insight into a key pathway (information search) through which implicit attitudes impact decision-making. Moreover, I empirically demonstrate important potential boundary conditions for understanding the influence of such attitudes on behavior.

Acknowledgements

While we don't get to see each other as much as we would like, please know that you are always in my thoughts and this dissertation is a testament to your love and support.

Finally, to my wife Marissa. I don't know where I would be in this world without you. This dissertation is as much yours as it is mine! We had been dating for less than a year when I asked you to come to New Jersey with me to start this crazy journey. Oh how far we have come! One marriage, two dogs, three apartments, and now one house later, and here we are. I can't believe how great these last six years have been and I am excited for what's to come. I can't imagine my life without you and I probably don't tell you this as much as I should, but you are truly the best thing to ever happen to me. Thank you for all of your encouragement, advice, patience, and general kindness that you exhibit on a daily basis. I love you.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

"Feelings can be present without awareness of them. The absence of conscious perception is no proof of the absence of mental activity."

Plotinus (c. 204-270 AD)

1.1 Motivation

Cognitive psychologists long ago demonstrated that the capacity to consciously process information is severely limited (Anderson, 1983, 2004; Norretranders, 1998). Lodge and Taber (2013) provide perhaps the most sobering assessment of such limited processing capacity by highlighting that approximately only two percent of information that we are exposed to (at any given time) enters our conscious awareness. With such serious limitations in conscious processing capacity, it is clear that decision-making is far more complex of a task than we often give it credit for. With this in mind, how is it that we are able to make decisions quickly and without expending immense amounts of energy?

The answer in part lies in how the human brain is hard-wired. The consensual view in cognitive psychology and neuroscience is that information processing and attitude formation operate via a dual process model of thinking and reasoning (De Houwer and Moors, 2005; Evans, 2003; Evans, 2010). Specifically, a distinction is made between unconscious or implicit "system 1" and conscious or explicit "system 2" level information processing. System 1 is characterized as spontaneous, fast, effortless processing, which operates below a conscious level of awareness. System 2, on the other hand, is

described as processing that is deliberative, effortful and self-aware.

To further define what is meant by an "implicit attitude" or system 1 processing, I will first provide a working definition of an attitude. Social psychologists typically characterize an attitude as a memory-based affective association between a category (or object) and its evaluation (Fazio et al. 1982, 1986; Bargh, 1997). With so-called "implicit" attitudes, the affective association between an object and the evaluation of that object is activated in a relatively automatic fashion (Bargh,1997; Fazio et al., 2007 Greenwald et al., 1998). By automatic, we generally mean that these affective evaluations are activated spontaneously, outside of one's full control or awareness, and expend few cognitive resources.[1] This is in contrast to system 2 processing or explicit attitudes, which are cognitively effortful, demand attention, and are presumed to be based on intentional search of the working memory (Lodge and Taber, 2013).

While this distinction between system 1 and system 2 level processing or "implicit" and "explicit" attitudes is a simplification and of course there is quite a bit of overlap in use between the two systems; there is a fairly extensive literature suggesting that system 1 processing or implicit attitudes enter the decision-stream before any cognitive considerations come consciously to mind (Zajonc, 1980; 2000). Empirical evidence shows that object-evaluation associations are formed within 200-300 milliseconds upon exposure to an object, which is well before cognitive considerations are activated (Burdein, Lodge, and Taber, 2006; Le Doux, 1994, 1996; Morris et al., 2003). This has led some to characterize such implicit processing as an "affect" heuristic (Slovic, 1999; Slovic et al., 2004, 2007; Sniderman et al., 1991), which represents a "quick and dirty" pathway for the brain to simplify thought and decision calculations that precedes conscious deliberation.

The claim that spontaneously activated attitudes, relatively inaccessible to conscious thought, potentially influence subsequent conscious deliberation and decision-making

---

[1] Although as Bargh (1997) and Djiksterhuis (2009) note, there is some debate on the extent to which implicit attitudes are truly "automatic." On some level this has to do with how we define conscious awareness. Dijksterhuis (2009) discussion on pre-conscious and post-conscious automaticity is relevant here.

has garnered significant attention across the social sciences. Research has primarily focused on the implications of such automatic processing and particularly the impact of implicit attitudes on social behavior (Ferguson and Porter, 2010; Hassin et al., 2005; Peugini et al., 2010). Numerous studies have used implicit attitudes to examine a variety of topics related to social preferences and judgments, including implicit racial bias (Jost, Banaji, & Nosek, 2004; Nosek, Banaji, & Greenwald, 2002), implicit gender stereotyping (Banaji & Hardin, 1996; Rudman & Glick, 2001), and implicit attitudes towards sexual orientation groups (Dasgupta & Rivera, 2006; Jost et al., 2004; Steffens & Buchner, 2003). More recently, there have been studies linking implicit attitudes to hiring decisions (Bertrand et al., 2005; Rooth, 2007; Rudman and Glick, 2001; Ziegert and Hanges, 2005), medical diagnoses (Green et al., 2007; Von Hippel et al., 2008), police behavior (Glaser and Knowles, 2008) and general interactions with outgroups (Dovidio et al., 1997, 2002; McConnell and Leibold, 2001; Rudman and Ashmore, 2007). In all of these studies, implicit attitudes were found to be a significant predictor of individual behavior, even when a respondent's explicit attitudes did not always have such an effect.

Starting with Lodge, Taber, and colleagues (Burdein, Lodge, and Taber, 2006; Lodge and Taber, 2000; Lodge and Taber, 2005), the use of implicit measures to examine political outcomes has grown in recent years. The aforementioned authors use a sequential priming procedure to develop their "hot cognition" hypothesis, which contends that political information processing operates via automatic affective evaluations. As the authors argue, these automatic affective evaluations activated in time t help drive biased deliberation and decision-making in time t +1. A second set of studies have examined the extent to which implicit attitudes toward social groups impacts individuals' policy preferences. Perez (2010), Knoll (2013), and Maholtra, Mo, and Margalit (2013) find that implicit racial attitudes are highly predictive of individuals attitudes toward U.S. immigration policy, even when controlling for explicit measures of racial bias and other types of confounds.

Implicit measures have also been used to examine candidate evaluations and voting

behavior. Todorov and colleagues (2005, 2007) explored how rapid judgments (often at the subconscious level) of facial images structure voters candidate impressions, which help to predict vote choice. Similarly, Carraro et al (2010) and Carraro and Castelli (2010) use implicit measures to explore the impact of negative campaign advertising on candidate evaluations. There have also been numerous studies on the impact of implicit racial bias on the 2008 U.S. presidential election. A number of studies (Finn and Glaser, 2010; Mo, 2011; Nosek et al., 2009 Pasek et al, 2010; Payne et al., 2010) report that implicit attitudes toward African Americans maintained a strong negative impact on the likelihood of voting for Barack Obama. Outside of the U.S. context, Acuri et al (2008) finds that implicit attitudes toward the two main Italian presidential candidates were highly predictive of vote choice, particularly among undecided voters. In a similar set of studies, Friese et al (2008, 2012) also find implicit attitudes to be highly predictive of future vote choice among decided and undecided voters in the U.S. and Germany.

In sum, a consensus is starting to emerge among those working in political psychology that much of political decision-making is at least partially influenced by implicit attitudes (Lodge and Taber, 2013). However, as with any new concept or measure, far more attention has been paid in trying to demonstrate both the construct and predictive validity of such attitudes, as opposed to a more serious examination of the pathways through which implicit attitudes may influence political behavior. This has led some to question whether the inclusion of implicit attitudes in models of political behavior adds significant explanatory power above and beyond the traditional measures of explicit or self-reported attitudes (Ditonto, Lau, and Sears, 2013; Kinder and Ryan, 2013; Segura and Valenzuela, 2010).

Additional criticism has been levied on the implicit attitude research program in political science on the grounds that little has been done to explicate the potential boundary conditions (or moderators) for understanding when implicit attitudes are likely to have their strongest and weakest effects on political outcomes (Huddy and

Feldman, 2009; Mitchell and Tetlock, 2009). Lodge and Taber (2013) echo such sentiment by highlighting the lack of knowledge pertaining to the processes underlying individual variation in the use of such implicit information processing. This has led some to argue that the new frontier of implicit attitude research will be to identify both the theoretical and methodological boundaries of implicit attitudes as it pertains to political behavior (Perez, 2013).

Hence, the goal of this dissertation project is to explore some of the pathways through which implicit attitudes impact voter decision-making. Specifically, this project will address two sets of research questions:

For my first set of research questions, I seek to unpack the often reported association between implicit attitudes and voter decision-making. One proposed but empirically under-developed claim is that implicit attitudes help to drive selective information processing, which in turn biases deliberative voter decision-making (Lodge and Taber, 2013). However, as Galdi and colleagues (2012) point out, previous studies focusing on implicit mechanisms of selective information processing have used as their primary predictor a set of explicit attitudes. A systematic study of this type using a direct measure of implicit attitudes (such as the IAT) has yet to be done. Hence, to assess the link between implicit attitudes and political information processing, I borrow from the work on behavioral decision theory and leverage the dynamic information board methodology (Lau and Redlawsk, 1997; 2006) to provide an empirical assessment of how implicit attitude strength impacts the amount, type, and duration of information accessed about political candidates in the context of a campaign.

Specifically for my experimental design, I construct a mock presidential primary campaign where subjects are presented a matrix of "learnable" information about fictional political candidates , which continuously flow down their computer screen for a limited length of time. Given the constant flow of information and the pre-defined time limit, subjects must make important tradeoffs as to both the amount and type of information accessed for each candidate. At the end of the campaign, subjects are

then asked for their vote choice followed by a series of candidate evaluation items (feeling thermometers and a memory recall task for each candidate). Finally, once this is complete, subjects are administered an Implicit Association Test (IAT) as of a measure of implicit attitude strength. Specifically, subjects are given an IAT to assess implicit racial prejudice, which is an area of inquiry that has gained significant traction in political science in recent years (Ditonto, Lau, and Sears, 2013; Kinder and Ryan, 2013; Mendelberg, 2001; Perez, 2010; Segura and Valenzuela, 2010). Accordingly, using the results from the racial bias IAT provides leverage in understanding both the broader impact of implicit attitudes on information processing in addition to the more narrow research question of how implicit racial bias impacts information search patterns in the context of a presidential campaign.

For my second set of research questions, I turn my attention to potential moderators of implicit attitudes' effect on voter decision-making. While a number of studies in political science have demonstrated the importance of implicit attitudes in a variety of political domains, much less is known about whether situational constraints and/or individual differences help to condition some of the demonstrated effects of implicit attitudes on political behavior. To explore such potential moderators, I use the same experimental design (discussed above) to conduct two additional studies of political information search and candidate evaluation but this time I account for two potentially important factors related to information processing: 1) the level of cognitive resources available to individuals (high or low); (2) an individuals' emotional state (positive or negative). The decision to examine these two features of the information environment draws on two distinct sets of literature, which contend that system 1 processing is more likely to be activated under certain situational contexts, including increased time pressure and cognitive load (Evans, 2003; Stanovich and West, 2000) as well as when individuals process information in a heightened emotional state (Desteno et al., 2004; Dasgupta et al., 2009). By accounting for these moderating factors, this experimental design allows us to assess important potential boundary conditions for understanding

when implicit attitudes are likely to have stronger (and weaker) effects on voter decision-making. To date, there has been no empirical assessment of such moderators in the context of political decision-making.

1.2 The specific political outcomes of interest in this dissertation

The focus of this dissertation is on the relationship between implicit attitudes and voter decision-making. As for the implicit attitude under examination, I chose to focus my attention on implicit racial attitudes toward African Americans. The social science literature on prejudice and politics is extensive. Numerous studies have demonstrated the effect of both explicit and implicit forms of racial bias on political behavior. Explicit prejudice has been shown to predict a variety of political behaviors including vote choice (Ditonto, Lau, and Sears, 2013; Kinder and Dale-Riddle, 2012; Kinder and Sanders, 1996; Kinder and Ryan, 2013; Mendleberg, 2001; Tessler and Sears, 2010) and opposition to policies such as affirmative action and welfare spending (Ditonto, Lau, and Sears, 2013; Gilens, 1999; Kinder and Sanders, 1996; Peffley and Hurwitz, 2007; Peffley et al., 1997; Sears, 1988; Sidanius et al., 1996).

Following suit, a growing number of studies in political science have started to explore the impact of implicit forms of prejudice on voting behavior (Greenwald et al., 2009; Mo, 2011; Perez, 2010; Payne et al., 2009; Pasek et al., 2009). Similar to previous studies on implicit racial prejudice, I use the 'black-white' IAT as my measure of implicit racial attitudes. The black-white IAT has been used in numerous published studies and has been shown to have rather high construct and predictive validity (see Bravo and Greco, 2013; Greenwald et al., 2001).

Turning now to the specific political outcomes of interest, this dissertation will shed new light on a number of outcomes related to voter decision-making. Fundamental to voter decision-making, I am first interested in how individuals search for information about political candidates. To the extent that the information gathering process

influences choice behavior, identifying an association between implicit attitudes and information search may go a long way to unpacking the effect of implicit attitudes on voter decision-making. Using the dynamic process tracing methodology, I am able to operationalize information search in a number of different ways.

Consider first the depth of information search. Here, I can simply examine the total amount of information accessed by a subject relative to all of the available information. A second set of key outcomes relates to the comparability of search across candidates. Specifically, I can examine the total number of pieces of information accessed for each candidate, which allows me to analyze whether subjects engage in different search patterns across candidates. Relatedly, I can also explore whether there are any differences in the type of information accessed for each candidate. By type of information, I mean information pertaining to either policy stances, personal information about the candidate, endorsements and polls. Finally, I can also compare the duration of time spent reading information about each candidate. In other words, in addition to tracking information search, I am also able to record the amount of time subjects spend reading any given piece of information (controlling for word count). This allows us to explore all of the same outcomes outlined above (i.e. depth of search, comparability across candidates, type of information accessed across candidates) but now with respect to reading time duration.

For a second set of outcome variables, I can shift focus from information search to candidate evaluation, choice behavior, and memory recall. To operationalize candidate evaluations, I include feeling thermometers (rated on a 0 to 100 scale) for each candidate. To examine choice behavior, I include a simple vote choice measure where subjects must "cast" their ballot in the mock primary election. Finally, to assess memory recall, subjects are administered a memory recall task where they are asked to write down as much information as they can about each candidate. This allows me to assess both the number of memories produced for each candidate as well as the accuracy of those memories.

1.3 Dissertation in a nutshell

To sum up, this dissertation is divided into seven chapters. In chapter 2, I further operationalize my working definition of implicit attitudes. This requires further elaboration on my initial characterization of implicit attitudes as object evaluation associations that are spontaneously activated, difficult to control, and operate below conscious level of awareness. I then shift attention to understanding how implicit attitudes enter into information processing. Specifically, I highlight the primacy of affect and how such automatic affective evaluations are hypothesized to influence subsequent conscious information processing and decision-making. Finally, I provide a selective review of recent articles that have incorporated implicit attitudes into the study of political outcomes. In my surveying of previous studies, I highlight important gaps in the existing literature and make note of where this dissertation project hopes to address some of these gaps.

In chapter 3, I provide a detailed overview of my experimental design. Unlike previous studies on implicit attitudes, this dissertation employs an experimental design that leverages the insights from behavioral decision theory to examine, at the most micro level, the relationship between implicit attitudes and political decision-making. Specifically, I use both the dynamic process tracing methodology (Lau and Redlawsk, 1997, 2006) and the IAT (Greenwald et al., 1998) to directly assess the link between implicit attitudes and voter decision-making. In addition to providing a detailed outline of these two procedures, chapter 3 also highlights both the cognitive load and emotional state manipulations. Chapter 3 finishes with a set of testable hypotheses derived from the experimental design.

In chapter 4, I use my experimental design to empirically assess the relationship between implicit racial attitudes and three sets of political outcomes: political information search, candidate evaluation, and vote choice. I find that implicit racial bias has a consistent negative effect on information search for the African American political candidate. Specifically, I report that implicit racial bias has a negative effect on both

the amount of opened information and reading time duration for the African American candidate. I report no such effect for the two white candidates in the primary campaign. Turning to candidate evaluation and vote choice, similar to other studies I show implicit racial bias to be negatively associated with both feeling thermometer scores and vote choice for the African American candidate. However, interestingly I find that this negative relationship loses significance once I account for information search patterns. Such findings suggest that one pathway through which implicit racial bias impacts voter decision-making is via biased information search.

In chapter 5, I turn my attention to potential moderators of implicit attitudes impact on voter decision-making. To do this, I introduce into my experimental design a manipulation of subjects' cognitive load (i.e. high vs. low) via the random assignment of a number memorization task that must be completed during the mock primary campaign. In terms of results, contrary to expectations, I fail to recover a significant IAT effect for any of my behavioral measures in the low cognitive capacity condition. Instead, I find that my explicit prejudice measure is a far stronger predictor of behavior when cognitive capacity is decreased. This is in stark contrast to my high cognitive capacity condition, where implicit racial biases significantly predicts most of my behavioral measures while explicit bias maintains no such effect. Upon first glance, these results appear to contradict the conjectures of dual process theory as to when implicit and explicit attitudes are likely to have their strongest and weakest effects.

However, as we detail in great length later on in chapter five, these unexpected results appear to be driven (at least in part) by an unintended consequence from my cognitive capacity manipulation. Specifically, I find that rather than overload subjects with information, this treatment condition appeared to unintentionally lead to subjects exhibiting higher levels of attention and engagement with the experimental tasks, thus crowding out much of the effect of implicit attitudes. Among other indicators, subjects in the treatment condition spent more time reading information about each of the political candidates and recalled more information than those subjects in the control. Reassessing the results with this in mind offers important insight into the potential

conditional effects of implicit racial bias. That increased attention and engagement with the decision task helped to crowd out the influence of implicit attitudes is consistent with the conjecture that explicit attitudes are better predictors of behavior when subjects are sufficiently motivated to engage in deliberative thought (Fazio and Towles-Schwen, 1997).

In chapter 6, I test a second potential moderating factor of implicit attitudes in the form of emotional state. In this chapter I examine subjects' current emotional state by using the PANAS-X scale for positive and negative affect (Watson and Clark, 1999). The PANAS-X scale has subjects assess (on a five-point scale) words and phrases that describe different feelings and emotions. From these ratings, summary scales are created to assess subjects' emotional state (positive vs. negative) and this is followed by statistical tests to identify whether there is a meaningful interaction between the salience of implicit prejudice and emotional state. Turning to my results, I find, consistent with affective intelligence, that the marginal effect of IAT (on information search) is indistinguishable from zero for those self-reporting to be in a negative affective state. In contrast, for those reporting to be in a positive affective state, the effect of IAT is negative and statistically significant. These findings provide initial empirical evidence to support the hypothesis that a positive emotional state facilitates the use of implicit attitudes while a negative one helps to inhibit it.

1.4   A note on some of the implications of this dissertation

The goal of this dissertation project is to further unpack how individuals make political decisions. From an information-processing and political-cognition perspective, I focus, at the most micro-level, on the individual-level psychological processes that help form political judgments and voter decision-making. By directly measuring individual differences in cognitive process (via a response time measure) as my main independent variable and using granular measures of decision-making (including information search,

candidate evaluations, and vote choice), I am in a position to make a number of important contributions to the burgeoning literature on unconscious or "implicit" forms of information processing and political decision-making.

First, my findings have important implications for much of the work in political psychology and public opinion on the impact of implicit attitudes on voter decision-making. Unlike previous studies, this dissertation project offers an experimental design capable of directly examining the impact of implicit attitudes on political information search and vote choice. By providing important insight into the relationship between implicit attitudes and information processing, this study unpacks how automatic associations can influence conscious beliefs and subsequent decision-making. In chapter 4, I found that implicit racial biases help predict information search on a variety of different dimensions. Such findings provide clear and convincing evidence to support the hot cognition hypothesis that automatic or implicit attitudes help drive selective information processing. Moreover, my results offer additional evidence suggesting that biased information processing is a key mediator for the relationship between implicit bias and political decision-making.

The second contribution made by this dissertation involves an in-depth examination of some potential moderators of implicit attitudes' impact on voter decision-making. By varying cognitive load and examining emotional state (respectively), this experimental design provides researchers traction on understanding some of the potential boundary conditions for the relationship between implicit attitudes and political decision-making. Specifically, in chapters 5 and 6, I reported important conditional effects for both implicit and explicit prejudice as a result of varying key features of the information environment and emotional state. The results from both these chapters provide initial empirical evidence to suggest that the impact of implicit attitudes, and specifically implicit racial prejudice, is at times context-dependent. Such findings also provided additional evidence in support of dual process models of cognition and behavior.

Finally, the results from this project help inform the growing debate on the significance of implicit attitudes as predictors of political behavior. While there has been recent push back by some scholars arguing that implicit attitudes add little additional explanatory power to models of voting behavior (Ditonto, Lau and Sears, 2013; Kinder and Ryan, 2013; Segura and Valenzuela, 2010), such findings may be underestimating the influence of such attitudes by focusing only on direct effects. The findings from this study suggest that implicit attitudes impact more subtle forms of political behavior, specifically how individuals search for information about politics. Such findings are consistent with previous studies that have found that implicit attitudes tend to have stronger effects on non-verbal forms of behavior (Dovidio et al., 1997).

Chapter 2

Literature Review

In this chapter, I further operationalize my definition of implicit attitudes. This requires further elaboration on my initial characterization of implicit attitudes as object evaluation associations that are spontaneously activated, difficult to control, and operate below conscious level of awareness. I then shift my attention to understanding how implicit attitudes enter into information processing. Specifically, I highlight the primacy of affect and how such automatic affective evaluations are hypothesized to influence subsequent conscious information processing and decision-making. Finally, I provide a selective review of recent articles that have incorporated implicit attitudes into the study of political outcomes. In my surveying of previous studies, I highlight important gaps in the existing literature and make note of where this dissertation project hopes to address some of these gaps.

2.1 The architecture of memory

A substantial literature in cognitive psychology has demonstrated that humans are limited capacity information processors (Simon, 1957, 1979; Miller, 1956; Taylor, 1981; Anderson, 1983). To better understand this, it is worth briefly reviewing the architecture of memory and how stimuli enter our conscious awareness. The dominant structural model of the mind is based on the classic architectural distinction between long-term memory (LTM) and working memory (WM) (Anderson, 1983; Atkinson and Shiffrin, 1968; Miller, 1957).

The best way to characterize memory activation and information processing is as a parallel process where the WM is acquiring and sending pieces of information to the

LTM and the LTM is, in turn, integrating this new information with old information already stored in memory and sending back to the WM this updated information. The LTM can be considered as the storage system for the brain that is organized by an associative network of inter-connected nodes. Cognitive psychologists generally consider storage in LTM as having practically an unlimited capacity. In contrast to the LTM, our WM has a severely limited capacity with only about 7 +/- 2 concepts able to coexist simultaneously (Miller, 1956).

Hence, when an individual is exposed to an object or stimulus, the WM processes (a limited amount of) this stimuli and thus begins the activation of related concepts already stored in the LTM (Collins and Loftus, 1975). The more frequently the individual is exposed to the same or similar objects or stimuli, the stronger the connections between the associative networks of nodes and thus the stronger the association between the new objects or stimuli and related concepts already stored in the LTM (Fazio, 2007). Consequently, these strong associations facilitate the retrieval of information from the LTM back to the WM. Thus, much of human's limited processing capacity stems from the inability to make strong associations between incoming stimuli or objects and the LTM. As such, most cognitive psychologists believe that only those concepts that are most activated in the LTM (and thus have the strongest associations) are able to be retrieved and enter back into the WM (Neisser, 1967; Ratcliff and McKoon, 1996).

2.2 How do implicit attitudes enter into information processing?

Given the structure of our memory, humans are severely limited with respect to the amount of conscious processing that can be performed. Such a task appears even more daunting when one considers the amount of available bits of information that we are exposed to at any given time.[1] With such constraints, how are we able to successfully

---

[1]According to Lodge and Taber (2013), the human brain processes one million bits of tactile information and 100,000 bits of auditory information, while at best we become aware of just five bites of tactile and thirty bits of auditory information per second.

deliberate and make decisions and not be overwhelmed by such a complex information environment? The answer, in part, stems from the number of heuristics, habits, and other mechanisms that greatly simplify thinking, reasoning, and thus decision-making (Abelson and Levi, 1985; Kahneman and Tversky, 1972, 1973, 1984; Tversky and Kahneman, 1973, 1974).

In this section, I will focus on one type of cognitive heuristic, automatic affective evaluations or "implicit attitudes," which have seen an explosion of interest over the last decade.[2] Implicit attitudes can be characterized as feelings (or affective evaluations) about an object or stimuli that are activated within a few milliseconds of exposure thus preceding conscious awareness or control. Given their speed and automatic nature, scholars generally consider implicit attitudes to be a "quick and dirty" pathway to the brain that allows for the extremely rapid processing of information, thus greatly simplifying decision calculations (Lodge and Taber, 2013). Characterizing implicit attitudes as a type of heuristic is consistent with dual process accounts of human reasoning, judgment, and decision-making.

Dual process theorists characterize information as being divided into two related but distinct systems (Evans, 2003, 2007, 2008; Kahneman, 2011; Stanovich, 1999, 2011; Stanovich and West, 2000). System 1 is generally labeled as the automatic or implicit system and is characterized by processing that occurs spontaneously, outside one's conscious awareness, generally difficult to control, and require little cognitive resources. In contrast, system 2 processing is often considered the explicit system and is described as processing that occurs at the level of conscious awareness by making use of the central working memory system. System 2 processing is characterized by limited capacity, a slower speed of operation, and requiring significant cognitive resources.[3]

---

[2]Heuristics can operate both at the conscious and subconscious level. For instance, Kahneman and Tversky (1972, 1973, 1984; Tversky and Kahneman, 1973,1974) highlight cognitive heuristics such as availability, representativeness, and anchoring that allow decision-makers to simplify complex judgments often outside of conscious awareness. Lau and Redlawsk (2006) focus on more "conscious" forms of heuristic use in the context of a political campaign, including party identification, campaign endorsements, and polls.

[3]While the system 1 and system 2 distinction is useful for descriptive purposes, it is important to

Scholars generally consider implicit attitudes or system 1 processing to enter the decision stream before any cognitive considerations come consciously to mind (Zajonc, 1980; 2000). In fact, a number of studies have demonstrated that affective evaluations are formed within 200 to 300 ms of exposure to an object or stimulus, which is before cognitive considerations can be activated (Burdein et al., 2006; Le Doux, 1994, 1996; Morris et al., 2003). This has led some to characterize implicit attitudes as an "affect heuristic" (Clore and Isbell, 2001; Slovic, 1999; Slovic et al., 2004, 2007), which signals an object's affective coloration to an individual almost immediately upon exposure thus greatly facilitating subsequent conscious processing.

With regards to how researchers think this process works in the brain, Lodge and Taber (2013) provide an excellent summary. Recall how information retrieval from the LTM to the WM operates. The accessibility of information in the WM is a function of the strength of the association between the evaluated object and concepts already stored in the LTM. According to implicit attitude scholars, as repeated evaluations are made, the affective tag linked to the object grows stronger. Thinking again about the primacy of affect, Lodge, Taber, and colleagues contend that upon immediate exposure of an object or stimuli, the affective tag is automatically activated, thereby signaling the object's affective coloration. This affective coloration then makes its way to conscious awareness via the activation of pathways linking the LTM to the WM. However, since pathway activation is a function of the strength of the object-evaluation association, only information that is congruent with the initial affective tag will be recalled into WM, and incongruent information will be inhibited, the implication being that initial automatic evaluations influence, at least in part, all subsequent encoding, retrieval, and comprehension of new information. [4]

---

note that such a distinction is an oversimplification. As Stanovich and Evans (2013) point out, these processing styles should be seen less as discrete types, but rather as systems that overlap in their use.

[4]This mechanism is further spelled out by Lodge and Taber's work on hot cognition (Lodge and Taber, 2000, 2005; Erisen, Lodge, and Taber, 2008), which will be discussed in more detail later on in this section.

2.2.1   The first wave of empirical research on implicit attitudes

In addition to the significant body of theoretical literature on the formation of implicit attitudes and how they are hypothesized to impact information processing, there is also a long history of empirical work in cognitive and social psychology underpinning such conjectures. The cognitive revolution that took place in the field of psychology in the 1970s and 1980s did much to establish the concept of automatic or implicit processing. Initial studies focused on demonstrating the existence of automatic memory activation via the examination of semantic associations. For instance, researchers studied whether the exposure to the word (or image of) "bee" instantaneously activated associated memories for the concept of "honey" or whether "boat" for "water" or "brick" for "building."

A variety of methods were developed to help demonstrate these automatic or implicit links. Perhaps most prominently, the sequential priming procedure (Rosch, 1975) had subjects categorize a "target" stimulus (a word or picture) after being presented with a "prime." Results from this procedure showed that subjects responses to the target is faster when it is preceded by a semantically associated prime (i.e. a prime that shares the same meaning as the target). Subsequent research using this procedure found that this result held true even when the prime was presented for a very short period or even when the prime was delivered subliminally (Neely, 1977). By the early 1990s, a fairly strong consensus emerged regarding the importance and ubiquity of automatic memory activation in information processing (Bargh, 1999; Ratcliff and Mckoon, 1988, 1994; Shiffrin and Schnedier, 1977).

Building on this work in cognitive psychology, social psychologists started to explore whether such automatic processing extended to attitude activation. Fazio et al (1986) adopted the sequential priming framework to examine what he called "automatic attitude activation." Fazio and colleagues started from the premise that an attitude represents a memory-based association between a category (or object) and it's evaluation. From this definition, they used the sequential priming procedure to explore

whether attitude activation was accelerated when the target and prime were of similar valence (or affect). Indeed, the authors found that subjects were quicker to evaluate positive and negative targets when primed with positive and negative words. Additionally, they found such effects even when subjects were instructed to ignore the primes, thus again suggesting the automatic and uncontrollable nature of such attitudes. In a followup study, Greenwald et al (1996) report similar results even when the primes were administered subliminally (under 200 ms).

This initial work in automatic attitude activation led others to explore whether such findings applied to social objects and stimuli. One area where implicit attitudes gained significant traction was in the study of socially sensitive topics, such as racial attitudes. In a series of important studies, Fazio et al (1995) explored the speed of evaluating positive and negative targets when primed with photos of black and white men's faces. They found that white subjects were faster to categorize positive (negative) targets when given the white (black) face prime. To explore whether such implicit attitudes correlated with explicit forms of behavior, Fazio and colleagues created an index of automatic racial preference for each experiment participant by calculating the difference in relative response time for the black-bad/white-good and black-good/white-bad pairs. They reported strong correlations between implicit racial attitudes and a self-reported measure of racial bias, the Modern Racism Scale (McConahay, 1986), when controlling for motivation to control prejudice. Moreover, they found that higher implicit attitude scores were associated with less friendly behavior toward black experimenters. Research of this sort led to a number of additional studies demonstrating the generally robust finding of both the implicit activation of prejudicial associations and their subsequent effect on a variety of explicit behaviors (Banaji and Greenwald, 1994; Blair and Banaji, 1996; Greenwald and Banaji, 1995; Dovidio et al., 1997).

Soon after this initial wave of research on implicit stereotyping, Greenwald, McGhee, and Schwartz (1998) developed what has become one of the most widely used measures

in social science research, the Implicit Association Test (IAT).[5] The IAT is a computer-based test that measures the response times associated with the rapid sorting of stimuli (e.g. black/white) into evaluative categories (e.g. bad/good). The basic intuition behind the IAT is that people perform sorting tasks with greater speed and accuracy when they can rely upon highly routinized cognitive associations than when task demands are in conflict with automatic mental links (Rudman, 2011). As a result, individuals sort stimuli more quickly and accurately when pairs of categories are closely associated (or "compatible"), and are slower to sort stimuli when pairs of categories are not closely associated (or incompatible). This difference in the speed and accuracy with which an individual reacts to paired categories (the "IAT effect") is considered an indicator of the association between categories used (or the stereotypes held).

The IAT has been widely used to gain empirical traction on the prevalence and consequences of negative stereotypes or prejudice towards members of a variety of groups and social categories, such as race (Jost, Banaji, & Nosek, 2004; Nosek, Banaji, & Greenwald, 2002), implicit gender stereotyping (Banaji & Hardin, 1996; Rudman & Glick, 2001) and gender and sexual orientation (Dasgupta & Rivera, 2006; Jost et al., 2004; Steffens & Buchner, 2003). Hundreds of articles using the IAT have been published in social psychology in the last two decades (see Jost et al., 2009 for an in-depth overview). Over time, IAT studies started to populate other disciplines, including organizational behavior, marketing, economics, public health, legal studies, the medical sciences, and political science. The spread of the IAT to other research areas has also contributed to a wider diversity of topics under study. Recent studies have examined how implicit attitudes may influence various forms of behavior such as hiring decisions (Bertrand et al., 2005; Rooth, 2010; Rudman and Glick, 2001; Ziegert and Hanges,

---

[5]Before I move forward with a review of more recent advances in the implicit attitude literature, it is important to emphasize an important distinction between two types of automaticity: pre-conscious and post-conscious automaticity. Our discussion of the sequential priming procedure and its variants is a measure of pre-conscious automaticity. In pre-conscious automaticity an object or stimuli is experienced below the threshold of conscious awareness. This is usually done via a prime with an exposure time under 300 ms. In contrast, with post-conscious automaticity, an object or stimuli enters conscious awareness, however, the individual is unaware of its influence on thoughts, feelings, or behaviors. The IAT is an example of a post-conscious automaticity measure.

2005), medical diagnoses (Green et al., 2007; Von Hippel et al., 2008), police behavior (Glaser and Knowles, 2008), product choice (Friese, 2006, 2008) and general interactions with members of out-groups (Dovidio et al., 1997, 2002; McConnell and Leibold, 2001; Rudman and Ashmore, 2007). In many of these studies, implicit attitudes had a significant impact on individual behavior, even when a respondents explicit attitudes did not.

2.3 The study of implicit attitudes in political science

The use of implicit attitude measures in political science has grown in recent years. This is, in part, evidenced by a recent special symposium in *Political Science and Politics* on implicit attitude research in political science (Hedrick & Ksiazkiewicz, 2013). There have also been a number of review essays highlighting work in political science using implicit measures in addition to discussing the potentially important implications of incorporating implicit attitudes into political science research (see Perez, 2013; Gawronski et al., 2013). Moreover, Lodge and Taber's (2013) recently published book detailing their JQP model of affect-driven political information processing has generated significant attention among political psychologists and the wider discipline alike.

In this section, I will briefly review some of the prominent studies that have used implicit measures to examine political outcomes. I start by highlighting some of the initial work by Lodge, Taber, and colleagues demonstrating that political objects, figures, groups, and symbols carry automatic affective tags and that implicit processing is a ubiquitous feature of political cognition. I then move on to reviewing a number of studies, which have attempted to correlate implicit attitudes with political behavior. Such studies span a variety of domains including policy issue positions, candidate evaluations, and vote choice. Finally, I will finish my review with a brief description of recent work that has tried to examine some of the pathways through which implicit attitudes impact political behavior and some potential moderating factors that may affect this relationship.

2.3.1   Implicit attitudes and political information processing

Lodge, Taber, and colleagues (2000, 2003, 2004, 2005, 2013) work on "hot cognition" played an important role in introducing implicit attitudes to political science. The "hot cognition" model argues that automatic affective evaluations largely condition information processing and political decision-making. Specifically, "hot cognition" says that all sociopolitical concepts (e.g. political leaders, groups, issues, etc.) are affect-laden (positively or negatively) and that this affect is linked directly to memory, where it is automatically activated upon exposure to a concept. To test their model, Lodge and Taber (2000) use an attitude priming test (Fazio et al., 1986), that records the speed (in milliseconds) with which individuals correctly identify the positive or negative valence of a word after being given a prime. As mentioned earlier, the idea behind the test is that if the prime and the target valences are congruent, then the response time will be faster. If the prime and target valences are incongruent then the response times will be slower.

In the Lodge and Taber experiments, political primes were used, which included people (Bush, Gore, C-store, etc.), groups (Republicans, Democrats, etc.), and issues (welfare, death penalty, etc.). The target words were affectively unambiguous (e.g. comedy, miracle, rainbow, toothache, demon, etc.). In terms of results, the authors report a highly significant interaction between the prime and target valences, thus providing support for the hypothesis that affect is triggered automatically with mere exposure to a political prime. This is further supported by Morris, Squire, Taber, and Lodge (2003), who find that their affective priming task (with political primes) led to systematic differences in brain wave patterns when subjects encountered affectively congruent and incongruent prime/target combinations. Evidence of affective priming was also found when individuals were exposed to in-group or out-group designations. In a series of experiments, Burdein, Lodge and Taber (2004) showed that group identifications also contained affective tags in that subjects were quicker (slower) to identify in-group (out-group) when targets were given a positive prime.

In sum, the Lodge, Taber, and colleagues research program on "hot cognition" provided much of the ground work for subsequent studies examining the impact of implicit attitudes on various forms of political behavior. By demonstrating that sociopolitical objects are affect-laden and can be automatically activated (or inhibited) by affective cues, these initial studies provide fertile ground for understanding the role of automatic affect in decision-making and reasoning as it pertains to politics. As I will discuss later on this section, evidence consistent with "hot cognition" has been subsequently used to help explain how selective information processing or "motivated reasoning" operates as a mechanism for biased decision-making.

2.3.2   Implicit attitudes and policy preferences

In addition to examining implicit attitudes in the context of information processing, researchers have also used implicit measures to explore socially sensitive political questions. One such topic where implicit attitudes have gained traction in political science is in the literature on attitudes toward immigration in the U.S. Researchers exploring the determinants of anti-immigration attitudes have used IATs as a measure of racial bias, which is often hypothesized to be an important predictor of immigration preferences, , however such attitudes are generally difficult to measure through self-reports.[6] In one study, Perez (2010) looks specifically at anti-Latino racial bias as a driver of mass-level immigration preferences in the U.S. To isolate the effect of Latino racial bias, Perez administers a Latino IAT to both a student and non-student sample. Perez finds that IAT scores maintain a significant negative impact on pro-immigration preferences even when controlling for a variety of socioeconomic indicators and explicit measures of bias including ethnocentrism.

---

[6]As noted by Perez (2013), among others, much of the appeal of implicit attitudes in the social sciences is that it provides researchers with a measure of socially sensitive attitudes that is largely immune to social desirability bias.

In a similar study, Knoll (2013) constructs an IAT to measure implicit nativist attitudes toward immigration. Specifically, Knoll is interested in whether there is a preference for a "traditional" version of American culture, which potentially maps onto anti-immigration sentiment in the U.S. In his IAT, Knoll pairs positive and negative words with images of symbols meant to represent traditional American values and stereotypes (e.g. apple pie, Statue of Liberty, etc.) and images representing "Latino-American" culture (e.g. a U.S. flag with a Mexico flag, a Hispanic grocery store in a downtown American urban area, etc). Interestingly, Knoll reports that most respondents indeed maintained an implicit bias toward nativist symbols and images. Additionally, Knoll uses the IAT score as a predictor of immigration policy attitudes and finds a significant negative effect even when including a variety of known confounds.

Shifting from implicit attitude directed toward Latinos, Malhotra et al., (2013) conduct an in-depth study of American attitudes toward H-1B visas (a plurality of which are granted to high- skilled Indian immigrants employed in the high technology sector) in order to test the impact of both economic and cultural threat as drivers of immigration attitudes. The authors employ a targeted sampling strategy of high technology counties in order to measure a specific set of skills (high technology) that are threatened by a particular type of immigrant (H-1B visa holders). To address potential racial bias while at the same time accounting for potential social desirability effects, Malhotra et al conduct an anti-Indian IAT in order to measure racial/ethnic tolerance absent self-editing. The authors split the sample by IAT scores above and below the median. They found that those with high implicit bias maintained lower support for Indian immigration (and immigration more generally) relative to those with low IAT scores.

2.3.3 Implicit attitudes, candidate evaluation, and vote choice

The area where implicit attitude research has seen the most growth in political science is in the study of candidate evaluation and vote choice. In one set of studies, researchers have explored how rapid judgments (often at the subconscious level) of facial images

structure voters' candidate impressions, which help to predict vote choice. Todorov and colleagues (2005, 2006) present subjects with pictures (facial images) of the winner and runner-up in congressional races and were asked to make a variety of trait judgments, including competence. Todorov et al found that competence ratings correctly predicted nearly 72% of election outcomes in the Senate and 67% in House races. In a follow up study, Ballew and Todorov (2007) manipulated mode of thought (i.e. conscious or unconscious thinking) by having two deliberation conditions. In the first condition, subjects were given unlimited time to assess the facial images and make a decision and they were also explicitly told to deliberate and make a good judgment. In the second condition, subjects were given a response deadline of two seconds. In line with studies of unconscious thought in social psychology (Dijksterhuis, 2004a, 2004b), Todorov et al report that the effect of competency ratings was significantly smaller in the first condition (conscious thought) as compared to the second condition (unconscious thought).

A second set of studies have used implicit measures to directly assess the link between implicit attitudes and candidate evaluation. For instance, Carraro et al (2010) and Carraro and Castelli (2010) use implicit measures to explore the impact of negative campaign advertising. While previous research has shown that negative campaigning has little net effect on the source candidate relative to the target (Lau, Sigelman, and Rovner, 2007), Carraro and colleagues report a different story when looking at implicit attitudes. Specifically, while they were able to replicate earlier findings on negative campaigning; they also found that negative campaigning led to reductions in implicit measures of competence and favorability for target candidates. This suggests that negative campaigning may indeed produce some of its intended effects, but that this operates at the unconscious level.

Finally, a significant number of studies have also examined whether implicit attitudes help predict individuals' vote choice. Nosek et al. (2009) examine whether negative implicit bias against African Americans help predict vote choice in the 2008 U.S. election. The authors sampled roughly 1,100 voting-age U.S. citizens one week prior to the election. They also include two self-report measures of racial bias including

feeling thermometer scores and a battery of symbolic racism items (Henry and Sears, 2002). The authors find that implicit and explicit racial bias both explain unique variance in vote choice. In a second study, Payne and associates (2010) use 2008 ANES Panel and Time Series data to examine the effect of implicit racial bias on vote choice. Similar to the results of Nosek et al (2009), Payne finds that implicit racial bias has a negative effect on voting for Obama net of explicit attitudes. The authors conclude that since both implicit and explicit prejudice each predict unique variance in vote choice, the decision to measure either in isolation would be underestimating the total effect of prejudice.

Using a slightly different research strategy, several studies have examined whether implicit racial attitudes led to a shift in support for John McCain and/or voter abstention. Pasek et al (2010) perform a series of multinomial logistic regressions as an alternative model specification to test whether implicit racial bias is associated with voting for Obama, voting for McCain, voting for a non-major candidate, or not voting at all. Again, in-line with previous studies, Pasek et al reports that implicit racial bias significantly shifts the predicted probability of voting for Obama downward and more towards voter abstention. Similarly, Finn and Glaser (2010) find that even when controlling for traditional determinants of vote choice (party ID, ideology) along with race and explicit racial bias, implicit racial bias is strong predictor of vote choice and even outperforms the explicit measure. They also find that implicit racial bias not only predicted a higher probability of voting for John McCain but also significantly increased the likelihood of not voting at all.

Attention has also been paid to how implicit attitudes may contribute to the prediction of political decisions relatively far in the future, specifically among undecided voters. Both Friese et al (2007) and Arcuri et al (2008) use IATs to assess implicit party evaluations several months prior to parliamentary elections in Germany and Italy respectively. In each study, the authors report that implicit attitude scores were highly predictive of vote choice, particularly among undecided voters. More recently, Friese et al (2012) assessed the usefulness of both implicit and explicit measures of candidate

evaluations for the prediction of future vote choice in the 2008 U.S. presidential election and the 2009 German Parliamentary elections respectively. The authors report that IAT measures of candidate and party preferences, obtained up to three months prior to the election, were better predictors of vote choice among both decided and undecided voters than the traditional explicit measures.

2.4   The 'downstream' effects of implicit attitudes

Researchers have amassed considerable evidence to suggest that implicit attitudes are potentially important predictors of a variety of political outcomes. Despite such evidence, it remains unclear the mechanism linking implicit attitudes to political decision-making. As Lodge and Taber (2013) point out, there are two ways that implicit processing can drive political behavior: (1) they may **directly** trigger a snap judgment or "gut reaction" that is entirely out of conscious awareness or thought (e.g. Todorov, 2005, 2007, 2010); (2) they may **indirectly** drive behavior through their impact on the process guiding political decision-making. While Todorov and colleagues have largely staked claim to the first option, most studies of implicit attitudes either explicitly or implicitly assumed the latter. However, the literature explicating the link between implicit attitudes and conscious thought processes is largely underdeveloped with only a handful of studies attempting to unpack the now the generally established association between implicit attitudes and various forms of political behavior.

Again, Lodge, Taber, and colleagues provide some of the initial research in this area. In Taber and Lodge (2006), the authors find evidence consistent with the theory that affect-driven processing produces selective information exposure. In their experimental design, subjects were given an information board and asked to search for information about two sets of political issues. Subjects were also asked to assess the strength of eight arguments (four pro and four con) for each policy issue. Taber and Lodge report that subjects were far more likely to access information and support arguments congruent with their pre-existing beliefs (as measured by an attitude battery administered prior

to the experiment). However, an important limitation to this study is that the authors do no directly measure subjects' automatic affective associations and instead rely solely on self-reported policy preferences as a proxy for automatic affective associations. In a more recent paper, Erisen, Lodge, and Taber (2013) address this critique and use the subliminal priming procedure to assess whether affective primes facilitate and/or inhibit subsequent conscious thought and deliberation. The authors find that subconsciously priming subjects with a positive or negative affective cue significantly altered their retrieval of information in the LTM and how they consciously rationalized arguments for or against a particular policy issue.

Perhaps, most relevant to this dissertation project is the recent work by Galdi and colleagues (2012). In their study, the authors investigate whether automatic associations trigger selective information search among those undecided about a particular policy. To test this, the authors administer a single category IAT that pairs positive and negative words with picture symbolizing the integration of Turkey into the EU. One week later subjects are given a selective exposure task where they are presented with pairs of newspaper headlines that suggested either a favorable or unfavorable article relating to the inclusion of Turkey into the EU. In line with their predictions, the authors find that, among undecided voters, selective exposure was significantly predicted subjects' IAT scores. The authors conclude that their findings point to a potential mediator for the link between implicit attitudes and individual decision-making.

A second set of studies have focused attention not necessarily on the mechanisms linking implicit attitudes to decision-making, but instead on potential moderators of this relationship. Dual process theory provides a useful starting point for thinking about potential moderators of implicit attitudes (Bargh, 1992; Bargh, 1999; De Houwer and Moors, 2005; Evans, 2003; Evans, 2010). As mentioned earlier, dual process theory contends that individuals possess two processing systems, one of which is automatic or implicit (generally referred to as system 1 processing) and the other of which is an effortful explicit attitude (generally referred to as system 2 processing). With regards to the conditions under which individuals are more likely to rely on system 1 rather

than system 2 processing, the general consensus in the dual process theory literature is that automatic attitudes have more of an effect in situations where cognitive resources are taxed (Evans, 2003; Payne, 2001; Stanovich and West, 2000). Similar work has been done by Fazio and colleagues (Fazio and Towles-Schwen, 1999; Fazio and Olson, 2003) with their conceptualization of the MODE model of behavior. The central tenet of the MODE model is that behavior is driven by controlled processes (i.e. system 2 processing) only when an individual is sufficiently motivated to engage in deliberative thought and has the cognitive capacity to do so.

Empirically, only a few studies have tested these propositions. Payne (2001) examines the strength of affective evaluations in light of time pressure. In his experimental design, subjects participated in a priming task where photographs of White and Black male faces were used as primes, which were then followed by targets that were either handguns or hand tools. As predicted, Payne found that subjects were more likely to incorrectly identify a tool as a gun when the target was primed with a Black face. To test for the effect of time pressure, subjects in the treatment group were randomly assigned a 500ms time deadline to respond to the target stimulus. Payne found that those given the time deadline were more likely to be impacted by the Black photograph prime.

In a second study, Friese et al (2008) examined the predictive validity of implicit attitudes pertaining to consumer behavior when varying subjects' cognitive capacity. Specifically, subjects were first administered both an implicit and explicit measure of their preference toward fruit and chocolate. They were then given a choice task where they were asked to choose five from a selection of ten fruit and ten chocolate items. To manipulate cognitive capacity, subjects are randomly assigned to one of two experimental conditions. In one condition, subjects must memorize a one-digit number during the choice task (i.e. high cognitive capacity). In the other condition, subjects must memorize an eight-digit number (i.e. low cognitive capacity) during the choice task. In terms of results, the authors report a significant interaction between cognitive capacity and their implicit measure (but no such interaction with the explicit measure)

with the IAT effect on choice behavior having a much larger effect in the low cognitive capacity condition. Hence, the authors conclude that the predictive validity of implicit measures does indeed increase when cognitive capacity is more limited.

Finally, in an entirely different domain, Desteno et al (2004) and Dasgupta et al (2009) examine whether emotional state can influence the strength of automatic evaluations. In Desteno et al (2004), the authors sort subjects into minimal in-group and out-groups. Subjects are then randomly assigned to a condition where one of three emotional states are induced: anger, sadness, or a neutral state. Finally, subjects are administered an IAT used to measure the strength of association between the in-group/out-group and positive/negative attributes. The authors hypothesize that anger should facilitate automatic evaluations of the in-group/out-group since anger is generally associated with system 1 processing and a greater readiness to use cognitive heuristics (Bodenhausen, Kramer, Susser, 1994; Bodenhausen, Sheppard, Kramer, 1994; Tiedens Linton, 2001; Marcus and Mackuen, 2001). In line with their prediction, the authors report a significant interaction between their anger manipulation and IAT score. They found no such interaction with their other emotional cues. In a similar study, Dasgupta et al (2009) use the same experimental design but this time test it for known out-groups. Hence, in this study they use Arab-Americans as the relevant out-group and again find that the strength of implicit prejudice against Arabs increased when subjects were primed with the anger emotional cue.

2.5   Outstanding issues

This brief and selective literature review points to a number of recent advancements in the IAT literature. Researchers have convincingly demonstrated that implicit attitudes can be important for political decision-making and do impact a variety of important political behaviors. With that said, there are number of important outstanding questions that need further exploration. For one, while recent work has made some headway on the mechanisms linking implicit attitudes to political decision-making, there is still much work to be done. Galdi et al (2012) offers the best attempt at addressing this

question by trying to empirically verify the link between implicit attitudes and selective information exposure, however a more careful analysis that leverages a variety of different ways to measure information search is needed. In this dissertation project, information search is operationalized through a number of different approaches including search patterns across alternatives, the type of information accessed, and the time spent processing information.

Moreover, the type of implicit attitude under examination may also be an important factor in determining how such attitudes enter (and impact) conscious thought. This may be particularly the case for socially sensitive topics such as racial and gender discrimination. Previous research has shown that implicit racial bias influences behavior in often subtle and more complex ways (Dovidio et al 1997) than other types of implicit attitudes. An important factor here is that individuals may be more motivated to control their implicit racial biases given its socially sensitive nature (Fazio and Towles-Schwen, 1997; Mo, 2011). This also may lead implicit racial biases being (perhaps more than other types of implicit attitudes) more likely to impact non-verbal behaviors than explicitly deliberative forms of behavior (Dovidio et al., 1997). Returning to the Galdi et al study on implicit attitudes and selective information exposure, the focus is on implicit attitudes toward Turkish integration into the European Union. While this is surely an important topic, its to be expected that such attitudes are likely not strongly crystallized nor are they as generalizable to a wide variety of behavior. Moreover, since such attitudes are not socially sensitive, their impact on conscious behavior may be inflated since such attitudes are not subject to potential self-censoring (Fazio and Towles-Schwen, 1997).

In addition to understanding the pathway through which implicit attitudes impact political decision-making, a second under-developed area of research pertains to potential moderators of implicit attitudes' effect. While there is some research on this topic in social psychology and marketing (Fazio and Towles-Schwen, 1997; Friese et al., 2009; Payne et al., 2001), this issue has largely been ignored by political scientists. This has led to a number of critiques of the implicit attitude research program in political

science. In their discussion on implicit racial bias, both Huddy and Feldman (2009) and Mitchell and Tetlock (2009) argue that little is known about interactive effects and under what conditions implicit attitudes are likely to maintain their strongest (and weakest) effects. Perez (2013) echoes this sentiment by arguing that one of the new frontiers of this research will be to identify both the theoretical and methodological boundaries of implicit attitudes, especially as it pertains to politics.

This has also led to some recent studies questioning whether implicit prejudice (or implicit attitudes more generally) has much of an impact on political judgment and decision-making above that of explicitly measured racial attitudes (Ditonto, Lau, Sears, 2013; Ryan and Kinder, 2013; Segura and Valenzuela, 2010). For instance, using 2008 ANES data, Segura and Valenzuela (2010) include both explicit and implicit measures of racial bias among white and Latino voters as predictors of vote choice and find that the implicit measure (in this case, the AMP) explained little additional variance above that of explicit measures of racial bias. Also using the 2008 ANES, Ditonto, Lau, and Sears (2013) report that the symbolic racism measure was a much stronger predictor of vote choice than their implicit measure (again, the AMP). In fact, the AMP was a significant predictor for only one outcome variable (an Obama feeling thermometer) among white respondents. Finally, Kinder and Ryan (2013) also use the 2008 ANES and use two measures of implicit racial attitudes (IAT and AMP) to examine a variety of political outcomes including vote choice, job performance, political participation, and a number of policy issue preferences (crime, welfare, immigration, etc.). For each outcome variable, they report very weak effects for both implicit measures while the effect for their explicit measure racial bias (the racial resentment scale) is consistently strong.

One conclusion that some may draw from these recent mixed findings is that implicit attitudes provide little added explanatory leverage in certain domains of political behavior. For instance, Kinder and Ryan (2013) argue that the political consequences of implicit measures of prejudice is generally limited, however such attitudes may be

more important in predicting "spontaneous" judgments and behaviors that occur without planning or deliberation. While this may be (in part) true, such a conclusion underestimates the potentially strong indirect effects of implicit attitudes on decision-making. There are two reasons to think this may be the case. First, as mentioned above, previous studies have shown that implicit attitudes tend to have stronger effects on non-verbal forms of behavior (Dovidio et al., 1997). Hence, previous studies may be underestimating the impact of implicit attitudes since they are assessing such effects at the end of the decision-stream (i.e. choice behavior). Focusing on the information acquisition stage of decision-making may may be more representative of the true effect of implicit attitudes. Secondly, such mixed findings on the effect of implicit attitudes may also reflect to the lack of attention paid to the conditional nature of implicit attitude's effect. In other words, as Kinder and Ryan (2013) note, it may be the case that implicit attitudes are better predictors of political behavior when decisions must be made under significant informational constraints.

To date, there has been little empirical inquiry into potential moderators of implicit attitudes as it relates to political decision-making and voter behavior. This is somewhat surprising given that informational constraints are often a key feature of political decision-making. As Lau and Redlawsk (2006) note, during an election individuals often encounter an overwhelming amount of information and must deal with cognitive limitations. In one of the few studies to explore potential moderators of implicit attitudes in political science, Kam (2007) shows that the impact of implicit bias against Latinos on preferences for an hispanic candidate is conditional on the absence of a party cue. In Kam's design, subjects are randomly assigned to a treatment condition where subjects receive information pertaining to the party identification of the candidates. Kam finds that implicit (and explicit) bias against Latinos is a significant predictor of vote choice only in control condition (i.e. no party cue). When party cues are given, the effect of implicit bias appears to disappear. Overall, this suggests that party identification may help individuals override their implicit bias.

In a more recent study, Mo (2011) attempts to model the interaction between implicit bias and partisan identification in the context of the 2008 U.S. presidential election. Mo argues that those who identify as strong partisans possess a stronger motivation to vote for the party candidate and as such should be able override any automatic predisposition against a candidate due to race. Similar to Kam (2007), Mo finds a statistically significant positive coefficient for the interaction between her implicit racial attitude measure and partisan strength. In her plot of the marginal effect of this interaction, Mo shows that for those with a strong party affiliation (either those that consider themselves a strong Democrat or strong Republican), the impact of implicit bias on vote choice is minimal at best. Conversely, for those who do not have strong party affiliation, implicit bias maintains a strong negative effect on the likelihood of voting for Obama.

While both Kam (2007) and Mo (2011) provide useful starting points for addressing potential moderators of implicit attitudes, it remains unclear the mechanism that is driving the overriding of implicit attitudes in this context, particularly since any number of correlates of partisanship may be responsible for the observed effect (i.e. partisan identification is not randomly assigned). Moreover, the absence or presence of partisan identification may actually reflect the extent of informational constraints present (which would be in support of the hypothesis presented in this study), however this is difficult to disentangle from other potential mechanisms that may arise from cueing party identification. In order to unpack the moderators of implicit attitudes, a cleaner experimental design is needed that allows us to better able isolate the potential mechanism(s) at work. The experimental design offered in this study provides such a research strategy. By manipulating the key features of the information environment, I am able to cleanly provide a direct test of potential moderators of implicit attitudes that is theoretically grounded in a significant body of literature in cognitive psychology on dual cognition.

2.6 Contributions of this dissertation

To address these gaps in the literature, this dissertation project offers an experimental design capable of directly examining the impact of implicit attitudes on political information search and vote choice. Previous studies that have explored the determinants of selective information processing have either neglected to measure implicit attitudes directly or have used crude measures of information search as their outcome variable. Using both the Implicit Association Test (Greenwald et al., 1998) and the Dynamic Process Tracing methodology (Lau and Redlawsk, 1997, 2006), I construct a mock presidential primary campaign to test whether implicit racial attitudes are associated with systematic differences in both information search patterns and voting behavior across subjects. By providing important insight into the source of biased information processing, this study unpacks how automatic associations can influence conscious beliefs and subsequently decision-making. Additionally, it provides a direct test of the "hot cognition" hypothesis (Burdein et al., 2006; Lodge Taber, 2005). To the extent that this dissertation project can provide a direct link between implicit attitudes and decision-making, the findings from this study may help to identify one source of motivated reasoning and help to predict how individuals' political preferences may change over time.

Secondly, in this study, I account for two features of the information environment meant to examine potential moderators of implicit attitudes' impact on voter decision-making. These manipulations are grounded in an extensive theoretical literature (Evans, 2003, 2007, 2008; Kahneman, 2011; Stanovich, 1999, 2011; Stanovich and West, 2000, 2013), which suggests that key features of the information environment play an important role in determining the extent to which implicit attitudes drive a particular decision and/or behavior. Specifically, I examine two features of the information environment : (1) cognitive capacity; (2) emotional state. The decision to manipulate cognitive capacity stems from an extensive theoretical literature, which suggests that cognitive capacity can facilitate or inhibit the use of implicit processing (Fazio and Towles-Schwen, 1997). Moreover, manipulating cognitive capacity has an important

theoretical basis in the context of political campaigns and elections. Limited cognitive resources or "information overload" is a common experience among most voters during the height of a political campaign (Lau and Redlawsk, 2006).

For the second set of moderating factors, I chose to examine emotional state (positive vs. negative) since automatic processing and affect are so intertwined. A number of studies have already demonstrated that specific emotional states lead to systematic differences in the conscious processing of information (Brader, 2005; Brader et al., 2008; Huddy and Feldman, 2007; Lerner, 2001; Marcus et al., 2000; Parker and Isbell, 2010). However, much less is known about how such emotions directly impact implicit processing. Specifically, while there is some evidence to suggest that emotional cues can help to activate certain implicit attitudes (Desteno et al., 2004; Dasgupta et al., 2009), there has been little work on whether emotional state acts to moderate the link between implicit attitudes and decision-making. In the political domain, this has important implications particularly with regards to how negative campaign ads influence their viewers. As mentioned earlier, Carraro and Castelli (2010) find that negative campaign ads had differential effects on individuals' implicit and explicit attitudes. The findings from this study on emotions can help further tease this out.

Chapter 3

Research Questions and Experimental Design

The main goal of this chapter is to highlight both the primary research questions of this dissertation project and to outline the research strategy to address these questions. In the first half of this chapter, I provide an overview of the three main issues that the dissertation seeks to address. Specifically, I examine first the question of whether implicit attitudes impact political information search and candidate evaluation. Next, I explore whether this relationship is moderated via changes to the information environment. To test this, I first manipulate individuals' cognitive capacity to see whether implicit attitudes have higher predictive validity (in line with previous work in dual process theory) when individuals' cognitive resources are taxed. Finally for a third set of research questions, I use the PANAS-X scale to measure emotional state in order to address a series of hypotheses regarding the link between emotions and implicit attitude activation. Specifically, I am interested in testing whether a positive emotional state facilitates implicit attitude activation and if negative emotional state helps to inhibit it.

In the second half of the chapter, I focus on the research strategy implemented to address these research questions. As mentioned earlier, I leverage a novel experimental design that incorporates both the Implicit Association Test (IAT) (Greenwald et al., 1998) and the dynamic processing tracing methodology (Lau and Redlawsk, 1997; 2006). Hence, for most of this section, time will be spent detailing how both the IAT and dynamic process tracing works along with the specific details of the experimental design. Attention will also be paid to the both my randomly assigned manipulation and the data collection strategy for my emotional state scale. This section will finish by outlining a series of testable hypotheses, which will serve as the basis for the analysis

and results section seen later in this dissertation.

3.1 Research questions

3.1.1 Do implicit attitudes impact political information search?

The general consensus among implicit attitude scholars as to how such attitudes impact behavior is that upon exposure to an object or stimuli, an automatic (positive or negative) evaluation is made, which colors any subsequent conscious deliberation (Lodge and Taber, 2013). For many, this indirect effect of implicit attitudes on conscious behavior is thought to be mediated by selective information processing. As Lodge, Taber, and colleagues' research program on motivated reasoning argues, initial automatic evaluations produce a type of "snowball" effect with subsequent conscious information processing operating via the retrieval (from LTM) of considerations congruent with the initial automatic affect generated. It is hypothesized that this search for and affinity to information consistent with one's initial affective evaluation leads to systematic biases in both information processing and the decision-making that follows.

Despite there being an extensive theoretical body of literature to bolster the claims made by Lodge, Taber, and colleagues, the empirical tests of implicit attitudes' impact on political information processing is largely under-developed. In general, previous studies focusing on potential implicit biases as a cause of selective information processing have used as their primary predictor a set of explicit attitudes (see Taber and Lodge, 2006 as an example). In fact, to this author's knowledge, only one study (Galdi et al., 2012) has directly tested the impact of implicit attitudes on information search. Moreover, given the research design of the Galdi et al (2012) study, there are strong reasons for why additional testing is needed (see Chapter 2 for more detail). Among other things, Galdi et al (2012) use a single crude measure of information search. More importantly, their implicit attitude measure (feelings toward Turkish integration into the E.U.) likely operates differently from more socially sensitive implicit attitudes such as racial bias (Fazio and Towles-Schwen, 1999; Fazio and Olson, 2003).

Hence, in this dissertation project, for my first set of research questions, I try to further unpack the link between implicit attitudes and political information processing. First, I explore whether I can identify a direct association between implicit attitude strength and information search patterns in the context of a mock presidential primary campaign. Using the IAT in conjunction with the dynamic information board, I can test whether IAT scores help predict, among other things, the comparability of search across candidates, the total amount of information accessed, the type of information accessed (e.g. policy stances, personal information, endorsements, etc.), and the total duration of time spent viewing information. In addition to information search, I also include a number of candidate evaluation and vote choice items in order to try to replicate previous findings demonstrating a link between implicit attitudes and voter decision-making. Using this novel experimental design, I am in a position to directly assess the extent to which implicit attitudes impact voter decision-making via biased information processing.

3.1.2 Do manipulations of the information environment condition the effect of implicit attitudes?

In the dual process literature, it is argued that automatic attitudes will have a larger impact on behavior in situations where the time and resources to consciously deliberate are limited (Evans, 2003; Payne, 2001; Stanovich and West, 2000; Payne, 2001; Westling, Mann, and Ward, 2006; Shiv and Fedorkhin, 1999). In contrast, explicit attitudes are likely to be a more significant predictor of behavior when constraints on information and decision-making are relaxed (Fazio and Towles-Schwen, 1999; Strack and Deutsch, 2004). Empirically, there is some evidence to suggest that implicit attitudes' predictive strength is moderated by cognitive capacity. In several studies (Payne et al., 2001; Gibson et al., 2008; Friese et al., 2009), researchers have manipulated subjects' cognitive capacity (via a cognitive load task) and found that implicit attitudes were a significant predictor of choice behavior only when cognitive resources were taxed.

However, to date there has been no such analysis in the context of implicit racial

bias. This may have potentially important implications given that the effect of implicit racial attitudes' may operate differently than other types of implicit attitudes (see Chapter 2). Moreover, none of these previous studies has examined choice behavior as it relates to political decision-making. There are several reasons to think that the interaction between implicit attitudes and cognitive capacity may be particularly relevant for political decision-making. It is well-established that voters rely on a number of simplifying heuristics to navigate the complex and over-abundant amount of information about candidates present in political campaigns (Lau and Redlawsk, 1997; 2006; Bartels, 2000; Lodge and Taber, 2000). This "information-overload" quality that typically characterizes political campaigns may also help to facilitate the activation and use of implicit attitudes to guide voter behavior. Thus, the impact of implicit attitudes may be particularly strong in context of a campaign where voters are trying to search for information and evaluate candidates under significant informational constraints.

Hence, for my second set of research questions, I am interested in empirically testing the conjecture that the impact of implicit attitudes on political information processing is conditional on information constraints. I do this is via a randomly assigned manipulation of subjects' cognitive capacity. First, I examine whether low cognitive capacity (i.e. high information constraints) facilitate the activation and use of implicit attitudes, while inhibiting the impact of explicit ones. Secondly, I test whether high cognitive capacity (i.e. low information constraints) help to inhibit the use of implicit attitudes while at the same time strengthening the effect of explicit attitudes. To address these questions, I leverage the same experimental design as mentioned earlier, with the exception that subjects in this study are randomly assigned to a manipulation where they are given a memorization task to perform during the mock primary campaign.

3.1.3   Does emotional state moderate the effect of implicit attitudes?

Previous research has shown that specific emotional states can have an important effect on information processing and decision-making (Damasio, 1994, 1996; Lerner and Keltner, 2001; Lerner and Tiedens, 2001). In political science, a number of studies have demonstrated that different emotions have distinct effects on political information search (Brader, 2005; Brader et al., 2008; Huddy and Feldman, 2007; Lerner, 2001; Marcus et al., 2000; Parker and Isbell, 2010). Much of this work stems from Marcus and colleagues' (1993, 2000) affective intelligence research program. These authors make a distinction between two affective sub-systems of the brain. The disposition system operates in familiar contexts and thus uses already established behavioral routines to deal with the external environment. The surveillance system, on the other hand, identifies stimuli that are unfamiliar and thus shifts attention away from pre-established routines. Marcus and Mackuen, along with others, further demonstrate that specific emotional states are more (and less) likely to trigger activation of the disposition and surveillance system respectively.

There are several reasons to think that the affective intelligence literature may be relevant to the study of implicit attitudes. For one, implicit attitudes are often defined as basic affective evaluations of objects (Smith and Nosek, 2011; Payne et al., 2005; Spence and Townsend, 2008). In other words, implicit attitudes reflect instant valence judgements (e.g. do you associate the target with good/bad, favorable/unfavorable, etc). Consistent with this definition, neuroscientists have shown that implicit (but not explicit) attitudes are associated with activity in the amygdala, which is the part of the brain that handles emotional responses to negative stimuli (Phelps et al. 2000; Cunningham, 2004).

In this study, my research questions center on whether distinct emotional states moderate the impact of implicit attitudes. To address this issue, I focus broadly on the relationship between positive and negative affect and implicit prejudice. Previous work (Marcus and Mackuen, 1993; 2000; Brader, 2005; Valentino, 2008) has shown

that negative emotions (such as fear/anxiety/hostility) trigger more active information processing and a shift away from easily accessible cognitive routines (i.e. activation of the surveillance system). In contrast, studies on positive emotions found that such affective states lead to increased heuristic use and more reliance on standard routines and stereotypes (i.e. the disposition system) (Albaracin and Kumkale, 2003; Isbell et al., 2006; Lerner and Keltner 2001).

Given that positive emotions such as enthusiasm tend to increase reliance on easily accessible routines, I test whether such an emotional cue also facilitates the use of implicit attitudes. Conversely, since negative emotions such as fear/anxiety is associated with a decrease in the reliance on heuristic use, I test whether such an emotional cue acts to inhibit the activation of implicit attitudes. To address these questions, I employ the same research strategy outlined earlier but now account for emotional state by examining subjects' score on the Positive and Negative Affect Scale (PANAS-X) (Watson and Clark, 1999), which was administered one week prior to subjects' lab sessions. The PANAS-X scale consists of twenty emotion items on which participants indicate their present feelings (1=very slightly or not at all, 5=extremely). Summary scales are then created to assess positive and negative affective states.

3.2 Design strategy

To address the three research questions highlighted above, I need a design strategy that allows for direct measurement of implicit attitudes, political information search, and candidate evaluation. Moreover, I must also specify a set of implicit attitudes that I am interested in studying. For this project, the decision was made to focus on implicit racial attitudes, specifically implicit bias against African Americans. This choice was made for a number reasons. First, there is already a robust literature on the relationship between implicit racial bias and voter decision-making (Nosek et al., 2009; Mo, 2011; Pasek et al., 2009; Payne et al., 2010; Finn and Glaser, 2010) and some debate on the magnitude of such effects (Ditonto, Lau and Sears, 2013; Kinder and Ryan, 2013; Segura and Valenzuela, 2010). Secondly, there is significant evidence to suggest

that given the socially sensitive nature of racial attitudes, implicit racial attitudes likely enter the decision-stream in more subtle and complex ways than other types of implicit attitudes. In particular, a number of studies have shown that individuals are more motivated (compared to other implicit attitudes) to suppress implicit racial bias (Fazio et al., 1995, Fazio and Towles-Schwen, 1997) and such attitudes are more likely to impact non-verbal forms of behavior (Dovidio et al., 1997). Hence, focusing on the impact of implicit attitudes on conscious behavior earlier in the decision stream (i.e. during information search) may provide a more accurate representation of the magnitude of implicit attitudes' effect. Third, outside of political science, the number of studies identifying the prevalence and use of implicit racial bias far exceeds the study of any other type of implicit attitude, thus suggesting the pervasiveness and importance of this implicit attitude for a variety of behaviors.

Turning now to the design strategy of this project, for the first part of the experiment I use the dynamic process tracing methodology (Lau and Redlawsk, 1997; 2006) to have subjects participate in a mock presidential primary election that includes two white and one African-American fictional candidates. During the campaign, subjects search for information (on a variety of different dimensions) about the candidates and at the conclusion, they are asked to evaluate each candidate and make a decision as to who to vote for. Once the mock campaign and election is complete, all subjects are administered the black-white race IAT (Project Implicit, 2012), which is used as a measure of implicit racial bias against African Americans. For the rest of this section, I will outline each component of the experiment in more detail.

3.2.1 Dynamic process tracing methodology

The dynamic process tracing methodology uses an information board to track all decisions made by subjects when learning about different alternatives (in this case different candidates). Information is presented to subjects on a computer screen and subjects must use the mouse to make selections as to what topics they are interested in learning more about. The information board contains six information boxes (at any one time),

which continuously move down the screen. Each box contains a label with a very short description of the nature of the information contained in the box and what candidate the information pertains to (e.g. "Lou Baker's Stance on Welfare"). Moreover, next to each information box is a photo of the candidate.

Unlike past work in behavioral decision theory that has used a similar (static) information board setup, the dynamic processing tracing methodology presents information in a "flowing" fashion.[1] Specifically, the six information boxes scroll continuously down the computer screen, with the box at the bottom disappearing and the new information box appearing at the top of the stack. Hence, every three seconds, subjects see the following: 1) a new information box appears; 2) all previously displayed information boxes move down one step; 3) one information box drops out of the information scroll area. The order in which information is presented varies randomly and as the subject proceeds through experiment, the computer continually records (1) the information boxes that are presented, (2) the order in which they are presented, (3) which boxes are selected by the subject, and (4) the time spent reading each opened information box. By clicking on the information boxes, subjects can learn more details about issue positions, candidate personalities, polls, endorsements, and background information about the candidates. When the subject finishes reading the information, they can return to the main screen where the information scroll continues to flow. Below is a screenshot of how information is displayed to subjects on the computer screen.

---

[1]This dynamic presentation of information has subjects face trade-offs with respect to what information to access since the available information is continuously changing. The idea here is that this setup more accurately reflects the 'real-world' campaign environment where the information environment is complex and individuals constantly face informational constraints.

The presidential primary campaign and candidates

In the mock U.S. presidential primary election subjects must choose to participate in either the Democratic or Republican primary. The choice of using a primary election serves to mitigate the overwhelming impact of party identification on information search and vote choice. In each primary, there are three candidates. The racial composition of the candidates is two white to one African-American.[2] In each primary, there are two moderate candidates and one highly partisan candidate. I assign policy positions to the candidates via independent coders ideology ratings of a wide range of policy stances.

---

[2]The decision to use three candidates stems from two concerns. First, the purpose of this study is to examine the impact of implicit racial bias (against African Americans). As such, I will need at least one African American candidate in the mock election. The issue is that if I only use two candidates (including one African American candidate), subjects may infer (at least partially) the nature of the experiment (i.e. that the study is in someway related to race). By using three candidates (including one African American candidate) for each primary election, the likelihood that subjects will make such an inference is much lower. Secondly, given that political campaigns tend to overwhelm voters with information, I wanted to do my best to replicate such an information environment while at the same time not making information so overly abundant that individuals are unable to retain any of the information that they were exposed to.

The composite average (i.e. the average of all issue position ratings for each candidate) was used to create an overall ideology score on a 1 to 7 scale. Below are the set of ideology scores for the three sets of candidates in each primary election.[3]

Candidate Ideology Scores for Both Primaries

Candidate	Ideology Avg.	Off Median
Liberal Candidate	1.94	-2.06
Moderate Candidate A (Democrat)	3.77	-.23
Moderate Candidate B (Democrat)	3.41	-.59
Moderate Candidate A (Republican)	4.59	.59
Moderate Candidate B (Republican)	4.65	.65
Conservative Candidate	6.09	2.09

By design, the two moderate candidates (in each primary election) differ in race (i.e. one white and one African American candidate). The highly partisan candidate is always white. To control for subtle differences between the two moderate candidates (in each primary), I randomly assign whether the African American candidate is "Moderate version A" or "Moderate version B." This washes out any potential differences in information search patterns and vote choice produced by the small differences in policy stances or personal characteristics between the two moderate candidates in each primary. In addition to blocking out differences in candidate ideology and personal characteristics, I also make some effort to control for the physical attributes of the candidates. All candidates are males, are of similar age (50 to 60), and are of similar physical attractiveness.[4]

---

[3]Please note that the ideology scores were adapted from Lau, Redlawsk, and Andersen, (2008).

[4]All potential candidate images were rated by independent coders for attractiveness. Only images with similar attractiveness scores were assigned to candidates.

3.2.2 How to measure implicit attitudes: the Implicit Association Test (IAT)

Perhaps the best known and most widely used procedure to measure implicit attitudes is the IAT, or Implicit Association Test (Greenwald et al., 1998). Numerous studies (Bar-Anan and Nosek 2012; Nosek, Greenwald and Banaji 2007; Egloff, Schwerdtfeger and Schmukle 2005; Schmukle and Egloff 2004, 2005) have shown that IAT scores are arguably the most reliable. The IAT is a computer-based measure involving the rapid sorting of stimuli (words and facial images) into two key paired categories. Subjects complete many trials, where each trial requires the sorting of a single stimulus item (usually words or faces that denote either "white" or "black"). Stimuli appear randomly and individually on the center of the computer screen and, as each stimuli appears, subjects must press the "E" or "I" key on a computer keyboard to quickly (and correctly) sort a given stimuli into the correct pairing ("white or good" and "black or bad" in some blocks of trials; and "white or bad" and "black or good" in other blocks of trials).

The logic behind the IAT is that individuals react more quickly to pairs of closely associated categories (Bravo and Greco, 2013). In contrast, individuals are slower to react to pairs of categories that are not closely associated.. This difference in reaction times (or the "IAT effect ") is considered an indicator of the association between the categories used. If one category contains evaluative attributes (e.g. words characterized as positive or negative) and the other category contains target stimuli (e.g. African Americans and Whites, or men and women, etc), then the difference score in reaction time is seen as an indicator of an individual's attitude toward that target category.

A good example of this procedure is the race (black-white) IAT used in this study, which attempts to capture implicit racial bias against African Americans.[5] The black/white IAT consists of seven blocks (see Table 3.1). In the first block (i.e. Demo block 1), respondents are asked to correctly categorize faces as being "White American" or "African

---

[5]For a demo version of this IAT, please see https://implicit.harvard.edu/implicit/demo.

American." The category labels (i.e. "White American" or "African American") for which respondents must sort stimuli are located in the top left and right corners of the screen. If the stimulus belongs to the category on the left, the respondent must hit the key on the left (in my experiment, the "E" key); if the stimulus belongs to the category on the right, the respondent must hit the key on the right (in my experiment, the "I" key). If the respondent incorrectly sorts the stimulus item, an "X" appears on the screen until the respondent presses the correct key. Demo 1 contains twenty trials (i.e. 20 sorting tasks).

In the second block (Demo block 2), respondents are given the same instructions, but now they must distinguish "Good" words (i.e. positive words) from "Bad" ones (i.e. negative words). Again, the category labels for which respondents must sort stimuli are located in the top left and right corners of the screen. Again, if the stimulus belongs to the category on the left, the respondent must select the key on the left (i.e. the "E" key); if the stimulus belongs to the category on the right, the respondent must select the key on the right (i.e. the "I" key). Demo block 1 also has a total of twenty trials.

The first two demonstration blocks are characterized as "single-categorization" sorting tasks since respondents are sorting stimuli on only one dimension (i.e. "White American or "African American"; "Good" or "Bad"). As such, they are not used to compute IAT scores. Instead, they are used as demonstrations to help train respondents to the nature of the sorting tasks. Prior to the administration of each of these first two blocks, subjects are given a screen of instructions that detail how to perform these sorting tasks.

The main blocks (for the scoring) of the IAT are compatible blocks 1 and 2 and incompatible blocks 1 and 2, which are the "double-categorization" tasks. These blocks use all four categories ("White American," "African American," "Good," and "Bad"). There are two versions of the double-categorization task. In the compatible version (i.e. compatible blocks 1 and 2), the category labels in each corner of the screen read "White American" *or* "Good" on one side and "African American" *or* "Bad" on the

other side. If the stimulus belongs to the category on the left, the respondent must hit the key on the left (in my experiment, the "E" key); if the stimulus belongs to the category on the right, the respondent must hit the key on the right (the "I" key).

In the second and incompatible version of the double-categorization task, (i.e. incompatible blocks 1 and 2), "White American" is now paired with "Bad" and "African American" is now paired with "Good." The category labels in each corner of the screen now read "White American" *or* "Bad" on one side and "African American" *or* "Good" on the other side. Again, if the stimulus belongs to the category on the left, the respondent must select the key on the left (in my experiment, the "E" key); if the stimulus belongs to the category on the right, the respondent must select the key on the right (in my experiment, the "I" key).

Each version of the double-categorization tasks consists of two blocks that are in successive order (i.e. compatible block 1 is always followed by compatible block 2; incompatible block 1 is always followed by incompatible block 2). Compatible block 1 (incompatible block 1) consists of twenty trials while compatible block 2 (incompatible block 2) consists of forty trials. Before the start of each block, respondents are shown a screen of directions that outline the nature of the sorting task in order to ensure that respondents perform the sorting task correctly.

For most IAT studies, the double-categorization tasks are administered in counterbalanced order. Approximately half of subjects are administered an IAT where the compatible blocks precede the incompatible blocks. The other half are administered the IAT with the incompatible blocks preceding the compatible ones. There is also one additional demonstration block (i.e. Demo block 3) that is administered between the two double-categorization tasks. The purpose of this block is to help respondents with the change in category labels that occurs when subjects move from the first double-categorization task to the second (for example a change in category labels where "White American" is first paired with "Good" and "African American is first paired with "Bad" to the new category labels where "African American is now paired with "Good" and

"White American" is now paired with "Bad.") To help ease this transition, respondents are administered a second single-categorization block (i.e. Demo block 3), which is same "White American" or "African American" sorting task in Demo block 1, but with switched response keys.[6] Demonstration block 3 consists of twenty trials.

A series of computer screenshots of the black/white IAT are provided below. For clarity purposes, I only display screenshots for the key sorting blocks in the IAT (i.e. the "double-categorization" tasks). The first set of screenshots are for the compatible blocks 1 & 2 (i.e. White paired with "Good"; Black paired with "Bad"). The next set of screenshots are incompatible blocks 1 & 2 (i.e. White paired with "Bad"; Black paired with "Good).

---

[6]As an example, if in Demo block 1, "White American" facial images belonged to the category on the right where respondents had to hit the "I" key and "African American" facial images belonged to the category on the left where respondents had to hit the "E" key; in Demo block 3 the position of the category labels and correct response keys would be switched.

Trial Blocks in the Black/White IAT

No. of trials	Block	Items assigned to left-key response	Items assigned to right-key response
20	Demo 1	African American images	European American images
20	Demo 2	Unpleasant words	Pleasant words
20	Compatible Block 1	Unpleasant + African American images	Pleasant + European Americans images
40	Compatible Block 2	Unpleasant + African American images	Pleasant + European Americans images
20	Demo 3	European American images	African American images
20	Incompatible Block 1	Unpleasant + European American images	Pleasant + African American images
40	Incompatible Block 2	Unpleasant + European American images	Pleasant + African American images

Note: The order of sorting tasks (i.e. "compatible" or "incompatible") is randomized in the administration of the IAT. However, it is always the case that Compatible block 1 (Incompatible block 1) precedes Compatible block 2 (Incompatible block 2).

Figure 3.1: IAT Compatible Task (i.e Compatible Blocks 1 and 2)
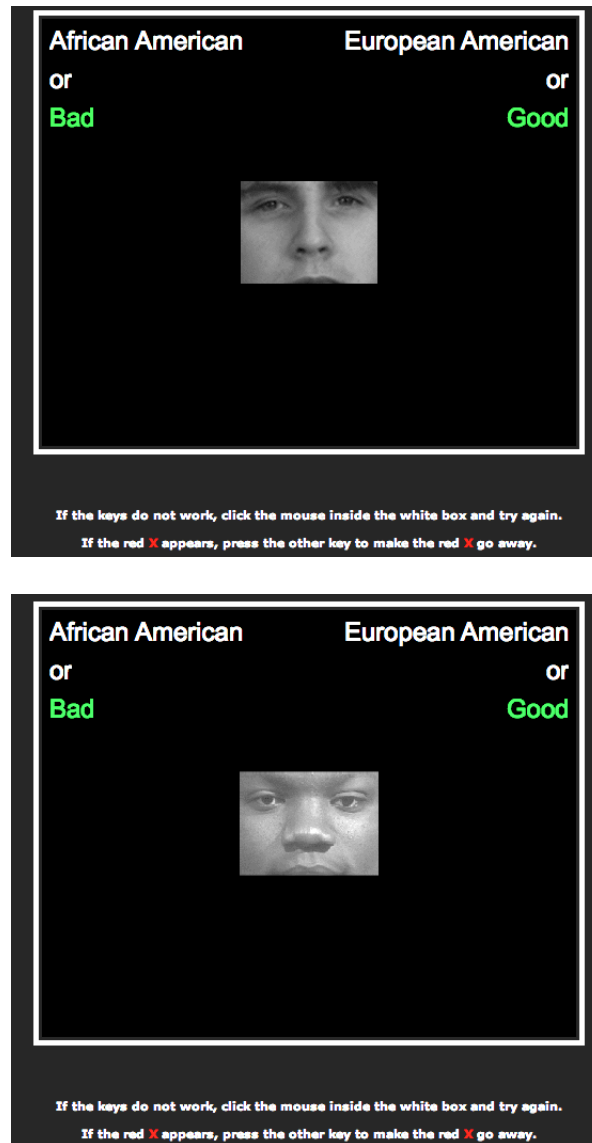
Figure 3.2: IAT Incompatible Task (i.e Incompatible Blocks 1 and 2)
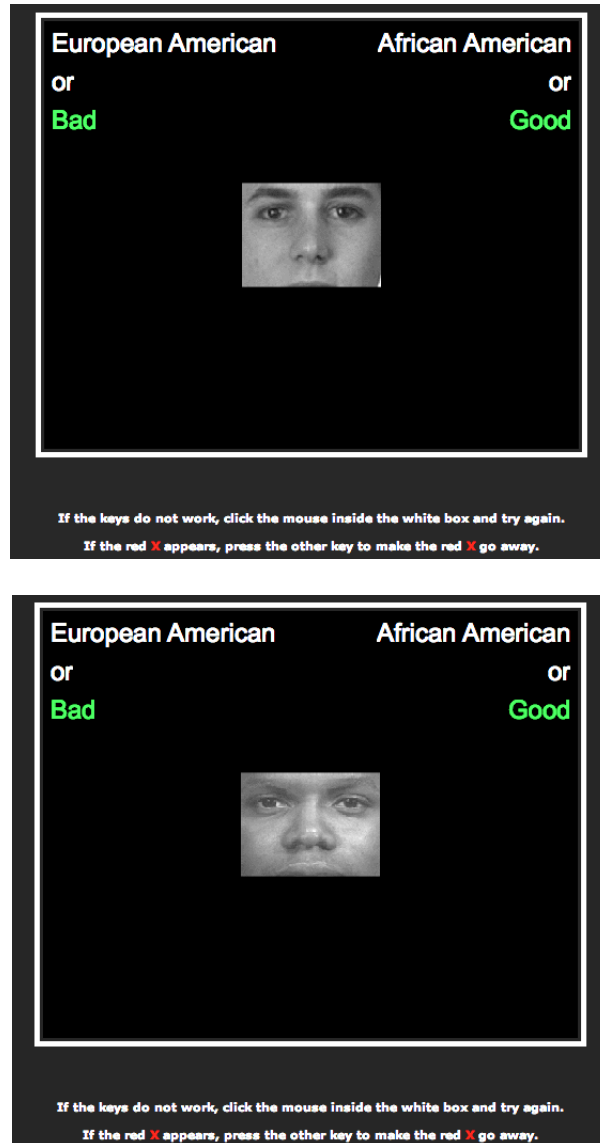
3.2.3   Why use the IAT?

I chose to use the IAT as my measure of implicit attitudes for this dissertation project for a number of reasons. First, the IAT is generally recognized as the most widely used measure of implicit attitudes (Jost et al., 2009; Lodge and Taber, 2013; Finn and Glaser, 2013). The reason for its popularity, in part, lies in its relatively high reliability compared to other implicit measures. A reliable measure is one that performs in a consistent and predictable fashion. There are different ways of assessing the reliability of a measure, including internal and temporal reliability. A scale is internally reliable (or consistent) to the extent that its items are highly intercorrelated. High inter-item correlations indicate that items are all measuring a shared underlying construct. The most commonly used measure of internal reliability is Cronbach's alpha, a coefficient that denotes the ratio between a measure's common and total variance. A second approach for assessing reliability focuses on a scale's temporal reliability. Test-retest correlations are usually used to assess temporal reliability.

In my discussion of the IAT's reliability, I highlight the race (black/white) IAT since that is the focus of this study.[7] Table 3.2 provides an overview of reliability studies for the black-white IAT. When assessing internal consistency, the race IAT scores appears to contain satisfactory measurement properties, i.e. Cronbach's alphas above 0.7, which is generally regarded as the minimum threshold for a "satisfactory" or "acceptable" measure. [8] Turning to temporal stability, the test-retest correlations for published race IAT studies tend to be in the 0.3 to 0.6 range for a variety of time intervals (as short as minutes and as long as two weeks). These are lower than what researchers typically expect from a stable measure.

---

[7]Previous work on the reliability of IAT scores has been developed in several substantive domains (race, gender, sexual orientation, age, and also a variety of "self concept" IATs looking at constructs like self-esteem). Since it is unclear whether the measurement properties of IAT scores vary by issue domain (and since my experimental study deals with the black/white IAT), I limit this discussion i to the reliability and measurement properties of the race IAT.

[8]Cronbach alpha scores are calculated by using all trials within each of the blocks.

Table 3.1: Studies Examining Test-Reliability of Black/White IAT

Study	N	Internal Consistency	r retest
Cunningham et al. (2001)	93	.88	.31 (14 days)
Bar-Anan and Nosek (2012)	100	.86	.40 (same session)
Dasgupta and Greenwald (2001)	48	NR	.65 (1 day)
Dasgupta and Greenwald (2000)	75	NR	.39 (same session)

One may wonder why is it that IAT is so widely used given these low test-retest correlations. However, this is not puzzling when one considers the reliability (internal and temporal) of IAT scores relative to alternative measures of implicit attitudes. IAT scores perform (relatively) well on measurement grounds when compared to the two other most commonly used measures of implicit attitudes. Most research suggests that IAT scores are more reliable than Affective Priming (Fazio et al. 1986; Fazio et al. 1995), and that IAT scores are at least as (and at times more) reliable as the Affective Misattribution Procedure (AMP) (Payne at al. 2005; Payne, 2009; Bar-Anan and Nosek, 2012), although there are few psychometric studies on the AMP at this point.

Let's first compare IAT and Affective Priming. In terms of internal consistency, Affective Priming generally yields very low split-half correlations, in the 0.2 to 0.3 range (see Fazio and Olson, 2003; Bosson et al., 2000; Perez, 2012).[9] The test-retest correlations for Affective Priming is also relatively low and subject to variability. Bosson et al. (2000) report a test-retest correlation of only .08 when the length between waves was roughly thirty one days. In contrast, Kawakami and Dovidio (2002) administered two waves of a supraliminal priming task in the same session, and found a higher (as expected, given that only minutes divided test 1 from test 2) test-retest correlation of 0.59.

Comparing IAT now to the Affective Misattribution Procedure, AMP scores have yielded split-half correlations in the 0.6 to 0.9 range (Payne et al., 2005; cfr. Bar-Anan and Nosek, 2012), which are similar, if on average slightly lower, than those of the IAT

---

[9]Split-half correlations are calculated by dividing each block into an equal number of trials and assessing the correlation between the two halves.

(see Bar-Anan and Nosek, 2012: esp. Table 2). Research on AMP's temporal reliability remains limited, but Bar-Anan and Nosek (2012) report test-retest correlations of the race AMP of 0.33 even when the time intervals between waves were very short (i.e. less than twenty-four hours).

3.2.4 Incentivizing the IAT

In this dissertation project, I use the novel approach of providing monetary incentives to subjects when administering the IAT. The motivation behind this design strategy is extensively detailed in Bravo and Greco (2013), however some detail on this is worth mentioning. Despite IAT's high reliability (relative to alternative implicit measures), an important literature in psychometrics has remained critical of IAT scores, given their susceptibility to measurement error and noise (see Blanton and Jaccard, 2006a, 2006b, 2008; also, see Arkes and Tetlock 2004, and Tetlock and Mitchell 2009). Like all latency-based measures, IAT scores are constructed using response times (to sorting tasks) and are measured in milliseconds. By using such a small unit of measurement, seemingly very small changes (say, a tenth of a second) in a latency can have a substantive effect on a person's IAT score. Such sensitivity can lead to low levels of reliability. As mentioned earlier, studies examining the measurement properties of the IAT have at times found reliability levels that are below the conventionally agreed upon standards for a "desirable" measure or scale, indicating that significant measurement error often accompanies IAT scores (Cunningham et al., 2001; Steffens and Buchner, 2003; Egloff et al., 2005).

These measurement issues are magnified by the fact that the IAT is fairly demanding on respondents in terms of attention and effort. Respondents are expected to remain closely engaged with the sorting tasks for several minutes, reacting to the stimuli presented, which need to be quickly (and accurately) sorted into the correct category via the pressing of one of two keyboard keys. Differential effort and motivation on part

of respondents may lead to considerable measurement error in IAT scores.[10] The heterogeneity induced by differential effort among IAT takers may lead to unknown forms of (systematic) measurement error that is not accounted for when used in statistical models.

In Bravo and Greco (2013), along with this dissertation project, I provide a performance-based monetary incentive to subjects administered the IAT. In Bravo and Greco (2013), they posit (and find) that a performance-based incentive helps reduce differences in effort level across and within subjects, thus addressing some of the measurement error surrounding the IAT and therefore increasing IAT's reliability. Empirically, the authors find that including a small performance-based incentive does increase the reliability of IAT scores, both in terms of internal reliability (i.e. Cronbach alphas and block correlations) and temporal reliability (i.e. test-retest correlations). Analyses of the latency data from Bravo and Greco's IAT study demonstrates that subjects in their treatment condition displayed response times with considerably lower variability and fewer extreme values across the blocks of sorting tasks that comprise the IAT. Furthermore, it is shown that this reduction in the variability of response times did not come at the expense of any accuracy losses (i.e. number of correctly sorted trials).

Hence, borrowing from Bravo and Greco (2013), in this experimental design all subjects are administered exactly the same standard race (black/white) IAT. I make no modifications whatsoever to this IAT. However, in each of my studies all subjects are informed that they are eligible to earn extra money depending on how quickly and accurately they complete their IAT. This is wholly in line with how the standard IAT is presented to test-takers, where subjects are actually asked to go *as fast and accurately as possible* when completing the IAT's sorting tasks.[11] Hence, all of the subjects in my experiment are asked to go as "fast and accurately" as they can when completing the IAT; however, in addition to these standard instructions, I also offer

---

[10]Krosnick and Lupia (2009) report that for the ANES sample, approximately 40% of IAT test-takers prematurely terminated their session due to frustration and fatigue with the sorting task.

[11]See the Appendix section 3.1 for exact wording of the instructions given to subjects

subjects a performance-based monetary incentive for doing so. Specifically, subjects are told that those in the top 20% (based upon both the speed and accuracy with which they complete the sorting tasks), will receive a $10 reward.[12]

3.3 Study procedures

In this section, I will briefly outline the procedures for each study in the dissertation project. It is important to note that in each study, the procedures are more or less the same. In each study there are two parts. Part 1 occurs approximately one week prior to the lab sessions with subjects receiving a pre-study questionnaire, which includes demographic items, political ideology, policy stances, and two measures of explicit racial prejudice.[13] In part 2, subjects arrive at the lab where they are given a short political knowledge quiz before starting the mock presidential primary campaign. The political knowledge quiz consists of eight items and covers a variety of topics including questions on the branches of government, identifying well-known political figures, and foreign policy.[14] Once subjects completed the mock primary campaign and voted in the election, they moved to the final part of the study, where I administered the black-white race IAT.

3.3.1 Study one

In study one, there are no manipulations to the information environment. Instead, the focus of this study is to simply explore the relationship between implicit racial attitudes and information search and candidate evaluation. Given the large body of literature on the primacy of automatic affective evaluations in influencing information search patterns and subsequent decision-making (Lodge and Taber, 2013), study one provides a direct test of this relationship by leveraging both the IAT and the dynamic

---

[12]Please see the appendix for the exact wording of the monetary incentive.

[13]I administered the demographic and explicit prejudice items one week prior to subjects' lab sessions so as to avoid the potentially biasing effect of administering the explicit and implicit prejudice items in the same session. For more details on pre-study questionnaire, please see the next section.

[14]See the appendix for a list of the questions.

process tracing methodology. If implicit attitudes impact political decision-making via differences in information processing, then I expect those with higher implicit racial bias to have different search patterns (and candidate evaluations) than those with low levels of such bias.

Turning now to the details of the primary campaign and election, as mentioned earlier there are three candidates in each primary. By design, in each primary there are two moderate candidates that differ in race (i.e. one white and one African American candidate) and one partisan candidate that is always white. The policy positions, personal characteristics, and endorsements for the two moderate candidates were randomized to ensure that subtle differences between the candidates did not influence information search patterns and candidate evaluation. Finally, all candidates were male, of similar age, and were of similar physical attractiveness.

As for the structure of the information environment, each candidate has forty unique information attributes that appear (twice) during the campaign. These forty information attributes include twenty issue stands, party affiliation, political ideology, seven items referencing personal background information, six personality descriptions, a picture, two performance evaluations, and two items pertaining to their campaign strategy. The campaign also includes endorsements from nine interest groups and the results from two different polls conducted at the early and late stages of the campaign. Candidate endorsements are randomly assigned (for the two moderate candidates) and the poll results always show each candidate to be neck-and-neck. In terms of timing, a new information box enters the screen every three seconds and the campaign in total lasts for approximately sixteen minutes.

Immediately after the campaign, subjects are asked to cast their ballot for the primary nomination. Images of each candidate are presented side-by-side to remind subjects of each candidate. After the vote choice item, subjects are administered feeling thermometers for each candidate. Lastly, upon completion of the candidate evaluations, subjects are given a memory recall task where they are asked to write down as

much information as they can remember about each candidate. [15]Once this is finished, subjects move onto the IAT portion of the experiment.

3.3.2 Study two

In study two, I am interested in assessing whether changes to the information environment moderate the effect of implicit attitudes. According to dual process theory, implicit attitudes should have their largest effects on individual decision-making when cognitive resources are limited such as in situations where time pressure is high and/or the available amount of information is overabundant. In order to empirically test these boundary conditions laid out by dual process theory, the treatment condition will manipulate the information environment by decreasing subjects' cognitive capacity.

Cognitive capacity or cognitive load refers to the use of working memory to process and learn information (i.e. to form schemas that link information processed in the working memory to the long-term memory) (Paas et al., 2004). Since working memory capacity is extremely limited (see Chapter 2), the amount of information that individuals can process simultaneously can either under-load or overload the amount of working memory available. If working memory is sufficiently overloaded, then information is not processed and thus will not be learned. The most frequently used method to manipulate cognitive load is to have subjects memorize strings of numbers while they are simultaneously performing another task (Swann et al., 1990; Shiv and Fedorikhin, 1999; Friese et al., 2008).

Following previous studies, subjects in study two are randomly assigned to a high cognitive load (or low cognitive capacity) treatment group. The high load treatment condition follows the same exact procedures as study one and does not differ in content (i.e. the candidates and the information available to subjects are exactly the same as study one). Rather, the only difference is the inclusion of a treatment, which seeks to manipulate subjects' cognitive load by asking them to memorize (and recall) sets

---

[15]It is important to note that the order of the candidates is randomized in the memory recall task.

of five-digit number strings while they participate in the mock presidential primary election.[16] Specifically, subjects are randomly assigned into one of three sub-conditions. In one sub-condition subjects are given two sets of number strings to memorize (and recall) during the campaign, but are not asked to memorize any number strings during the candidate evaluation/vote choice phase.[17] In a second sub-condition, subjects are only given one number string to memorize (and recall) at the onset of the candidate evaluation/vote choice phase.[18] Finally, in a third sub-condition, subjects are given the memorization task both during the campaign and during the evaluation period.[19] In order to assess whether the cognitive resources manipulation had its intended effect, I can use as a manipulation check the memory recall task given to subjects at the end of the election. Under the assumption that high cognitive load impairs individuals' ability to learn new information, I expect that memory recall should be greater in the low load condition relative to the high load condition.

The decision to include three sub-conditions for the cognitive load manipulation stems from the fact that it is not entirely clear how the cognitive load manipulation will operate. Previous studies that have used the memorization task to manipulate cognitive load typically include a choice task that takes no longer than a minute. Since the choice task is quite short, previous studies only use one set of number strings to manipulate cognitive load. In this study, there are two sets of choice tasks: one during the primary campaign and another when subjects must vote for a candidate. During the primary, the choice tasks consists of what information to look at for each candidate. Unlike previous studies, this choice task is quite long (approximately sixteen minutes) and continuous (i.e. choices are constantly being made). Hence, it is unclear as to whether the memorization of one number string during the entirety of the mock

---

[16]For the exact wording of this manipulation, please see the Appendix section 5.2.

[17]Specifically, one number string is given right before the start of the primary and the recall question is approximately seven minutes later. A second number string is given in the ninth minute of the primary and the recall question is again asked seven minutes later.

[18]The number string is administered immediately after the information search and the recall question is asked after the vote choice, feeling thermometers, and memory recall tasks are finished.

[19]In this condition, a total of three number strings are administered. Two during the information search (approximately seven minutes apart) and one at onset of the candidate evaluation stage

campaign (and election) is the most effective way to manipulate cognitive load. To account for this uncertainty, I use three different manipulations that vary the number of memorization tasks (and time between recall). This will allow us to explore whether differences exist on account of the number of memorization tasks given to subjects. If not, I can simply aggregate this data into condition. However, if there are significant differences, then I will treat each sub-condition as a separate manipulation.

3.3.3 Study three

In study three, I am interested in exploring a second set of moderators of implicit attitudes related to individual level differences. Here I explore whether emotional state moderates the impact of implicit attitudes on voter decision-making. A significant literature in political science has shown that emotions can trigger distinct forms of cognitive processing, which have important effects on decision-making. There are strong reasons to believe that implicit attitudes are in some way related to emotional state. Phelps (2000) examined the relationship between IAT (as measured by the "black-white" IAT) scores and amygdala activation, a region of the brain involved in emotional responses to threat, such as fear. Phelps reports a strong positive association between implicit racial bias and amygdala activation. Interestingly, Phelps fails to find such an association with explicit racial bias (as measured by the Modern Racism Scale). That implicit attitudes appear to be associated with the region of the brain involved in emotional response may have important implications for understanding potential moderators of such attitudes. Moreover, these results suggest that implicit attitudes may be in some way connected to the surveillance system in the brain as highlighted by affective intelligence theory (Marcus et al., 2000; Brader, 2005, 2008).

To further explore the relationship between implicit attitudes and emotions, I focus attention on subjects' PANAS-X score to assess emotional state. The PANAS-X scale asks subjects to assess (using a five-point scale) words and phrases that describe

different feelings and emotions. For instance, words such as afraid, scared, and nervous characterize fear/anxiety (i.e. negative affective state), while words like cheerful/excited/delighted characterize happiness/joviality (i.e. positive affective state). For each word, subjects are asked to "indicate the extent you have felt this way during the past few weeks."[20] Hence, for the analysis I create a summary scale for positive and negative affective state (respectively). I then examine possible interactions between these emotional states and IAT scores.

3.3.4 A note on study administration

To complete the three studies, two separate sets of lab sessions were conducted. These sessions were conducted in consecutive weeks, with subjects randomly assigned to attend a lab session either in week one or week two. For the lab sessions conducted in week one, subjects were administered the experimental design without any manipulations to the information environment. For the lab sessions conducted in week two, subjects were given the experimental design that included the manipulation to cognitive load.

The data collected during these two sets of lab sessions were used for each of the three studies. The results for study one are based on the data collected from the lab sessions conducted in week one (i.e. the experimental design with no manipulations). For study two, data collected in both week one and week two was used during analysis. Specifically, I used the data collected during week one (i.e. the experimental design with no manipulations) to serve as the "control condition" and the data collected during week two (i.e. the experimental design with the cognitive load manipulation) as the "treatment" condition. Finally for study three, I used for analysis the data collected from study one. This is appropriate for two reasons. First, in both sets of lab sessions, subjects were administered the PANAS-X survey, which allowed for the collection of data on emotional state. Second, since study one contained no manipulations to the information environment, it served as a more suitable choice for cleanly assessing the relationship between emotional state and information search.

---

[20]Please see the Appendix section 5.4 for the actual wording of this manipulation check.

3.3.5   Control variables

In order to cleanly assess the effect of implicit prejudice on both information search and candidate evaluation, I will need to account for a number of potential confounds. For one set of controls, I will simply account for demographic features of my sample, including party identification, gender, and race. In addition to demographics, I will also account for subjects explicit racial attitudes. This will allows us to examine the effect of implicit attitudes net of explicit ones. I measure explicit racial attitudes in a variety of ways. Specifically, subjects are administered both the symbolic racism scale (Kinder and Sears, 1981; Henry and Sears, 2002) and Kam and Kinder's (2008) ethnocentrism scale[21] In addition to explicit racial attitudes, I also include in the questionnaire numerous measures pertaining to interest and activity in politics.[22] Moreover, I include a measure of political sophistication, which is a scale of several survey items testing political knowledge.[23] Finally, I also include a baseline political preferences survey, which has respondents evaluate a number of public policy issues.[24]

As mentioned earlier, in order to avoid inducing potential biasedness by administering my explicit racial attitude items before (or after) the lab portion of the experiment, all data on these control variables is collected in a separate study session five to seven days prior to subjects' scheduled lab session. By spacing these survey items at a sufficient distance from both the dynamic process tracing and IAT portions of the experiment, I significantly reduce the risk of order effects and the possibility of inadvertently informing subjects about the purpose of the study.

---

[21]For the wording of these items please see Appendix sections 5.5 and 5.6.

[22]See the Appendix section 5.7 for exact wording of these items.

[23]See Appendix section 5.8 for exacting wording of these items.

[24]See Appendix section 5.9 for exact wording of each item.

3.4 Proposed analysis and hypotheses

For the analysis portion of this project, I am in a position to examine both information search and candidate evaluation as outcome variables. Borrowing from Lau and Redlawsk (2006, 2013), I can operationalize information search in a number of different ways. Most important for my purposes, I am interested in the comparability of search across alternatives (i.e. candidates). One way to measure this is to simply count (and compare) the total number of pieces of information accessed by a subject across each candidate. However, one potential issue with such a measure is that there may be considerable variance in terms of subjects' depth of search (i.e. how much total information they access). Such differences may skew my comparability measure. To account for this, I calculate the information search variable as a ratio between the total number of opened items for a given candidate to the total number of opened items across all candidates. As mentioned earlier, the dynamic information board also records (in seconds) the amount of time spent reading each opened information box. Hence, similar to the opened items variable, I can calculate the same ratio for the time spent reading information about a particular candidate.[25]

For candidate evaluation, the measures are fairly straightforward. First I can examine who each subject voted for in their primary. Secondly, I can compare feeling thermometer scores for each candidate. Finally, I can analyze the results from the memory recall task and assess both the number of memories produced for each candidate and the accuracy of those memories.

Given the setup of the experimental design and some of the potential outcome variables that can be investigated, let's focus on some of the hypotheses generated from this work. First, I should be able to say something about the relationship between implicit attitudes and information processing. More specifically, the first set of hypotheses are

---

[25]It is important to point out that the way comparability of search is defined for this study differs from that typically used in the decision-making literature (Lau and Redlawsk, 2006). In this study, since there is an attempt to "control" for individual differences in depth of search, my search variable is better characterized as a comparability of depth of search.

as follows:

H1a: All else equal, negative implicit racial bias toward African Americans will be negatively associated with information search for the African American political candidate.

H1b: All else equal, negative implicit racial bias toward African Americans will be negatively associated with the time spent reading information about the African American political candidate.

These hypotheses serves as one of the foundations for this project. I can address these questions by examining the association between IAT scores and a number of information search variables focusing on the comparability of depth of search. If I fail to reject these hypotheses, then I am on solid ground to make the claim that implicit attitudes influence voter decision-making, at least in part, through it's impact on information processing. This provides added insight on the potential mechanism linking implicit attitudes to voting behavior. Moreover, failing to reject these hypotheses also provides ample support for the "hot cognition" model (Lodge and Taber, 2000, 2005). Finally, it is important to note that I can assess the impact of implicit racial attitudes while also accounting for the effects of explicit prejudice. This allows us to examine how both implicit and explicit forms of racial prejudice impact information search.

H2a: All else equal, negative implicit racial bias toward African Americans will be negatively associated with candidate evaluation and the likelihood to vote for the African American candidate.

This hypothesis is in-line with previous literature on implicit racial bias, candidate evaluation, and vote choice, which has generally reported a significant relationship between IAT score and voter decision-making (Nosek et al., 2009; Pasek et al., 2009; Mo, 2011). Moreover, it represents the sort of final outcomes that I seek to explain via implicit attitudes' role in producing selective information search. Hence, I can explore whether IAT has a negative effect on feeling thermometer score, vote choice, and the number of unique pieces of information recalled for the African American candidate.[26]

---

[26]Again, in assessing this hypothesis, I will also account for explicit racial prejudice in order to examine how both implicit and explicit bias impacts candidate evaluation and vote choice.

H2b: The effect of implicit racial bias on candidate evaluation and vote choice will operate primarily through political information search

Although I am limited in the extent to which I can say something about causal mediation in this study[27], I can analyze the relationship between implicit prejudice and candidate evaluation/vote choice once I include the information search variables into the model specification. Specifically, I can examine whether the predictive strength of implicit prejudice on candidate evaluation and vote choice is attenuated once I account for information search. Such a finding would suggest a strong indirect effect for implicit prejudice via candidate information search.

H3: There will be a significant interaction between implicit racial attitudes and political sophistication

One potential interaction effect not related to my experimental manipulations involves the impact of political sophistication on implicit attitude strength. Holding the information environment constant, political sophistication may also have a moderating effect on implicit attitudes. One hypothesis that can be tested using this design is whether political sophisticates are able to override their implicit attitude. Hence, I can observe whether the effect of implicit attitudes is weaker for political sophisticates relative to non-sophisticates. An alternative hypothesis, and one advocated by Lodge and Taber (2005), is that affective evaluations will be stronger for political sophisticates since these individuals have been more frequently exposed to political primes and thus the strength of the affective evaluations should be stronger. In this project, I can test this hypothesis by simply interacting IAT scores with the results of my political knowledge quiz.

H4a: The effect of implicit racial attitudes on information search and candidate evaluation will be stronger (weaker) when there is increased (decreased) cognitive load.

H4b: The effect of explicit racial attitudes on information search and candidate evaluation will be weaker (stronger) when there is increased (decreased) cognitive load.

---

[27]As Imai et al (2013) point out, to effectively conduct causal mediation analysis (using the potential outcomes framework), the proposed mediator must be randomly assigned.

These hypotheses are in-line with some of the arguments made in dual process theory. As mentioned earlier in this paper, automatic attitudes are generally believed to be more salient given certain informational constraints. As such, the strength of implicit and explicit attitudes as predictors, should vary in-line with the structure of the information environment. To examine these hypotheses I can analyze two sets of outcome variables. I can first focus on information processing variables and see whether the impact of implicit and explicit attitudes on 1) depth of search and 2) comparability of search is conditional on my manipulation of cognitive resources.

Secondly, I can also examine whether the strength of implicit attitudes as a predictor of vote choice and feeling thermometer scores is also conditional on the information environment. Here, I can analyze simply whether the effect of IAT score on the likelihood of voting for the African American candidate in fact varies on account of my cognitive manipulation. If this is indeed the case, it presents several important implications for previous literature demonstrating an association between IAT and candidate evaluation. Perhaps most importantly, such a finding would suggest that there is heterogeneity in IAT effect size as a result of variations in the information environment. Moreover, by not considering the nature of the information environment, scholars may be routinely underestimating (or in some cases overestimating) the effect of implicit attitudes.

H5a: The effect of implicit racial attitudes on information search and candidate evaluation will be weaker for those in a negative emotional state.

H5b: The effect of implicit racial attitudes on information search and candidate evaluation will be stronger for those in a positive emotional state.

H5c: The effect of explicit racial attitudes on information search and candidate evaluation will be weaker for those in a positive emotional state and stronger for those in a negative emotional state.

These three hypotheses follow closely from the literature on emotions and decision-making. Higher levels of negative affect should activate the surveillance system of the brain since levels of threat will be increased. Accordingly, individuals in a negative emotional state should engage in more deliberative cognitive processing and information-seeking. This increase in deliberative processing should crowd-out the effects of automatic processing. Hence, for these individuals I should see the number of information attributes accessed increase. More importantly, the association between IAT and both information search and vote choice should be attenuated while the impact of explicit attitudes is likely to be strengthened.

In contrast, for those in a positive emotional state, the effect of automatic attitudes should be stronger (relative to the anxiety appeal) since these subjects should rely primarily on their disposition system. Since the brain is not reacting to threat, subjects will process information by relying, in part, on their automatic predispositions. Hence, for these individuals I expect IAT scores to more predictive (than explicit attitude measures) of both information search and candidate evaluation. The results from this study have important implications for the examination of implicit attitudes in the context of voter decision-making. In particular, it may shed light on whether the activation of automatic attitudes is susceptible to framing effects based on emotional appeals. Such results are important for the literature on the effects of negative advertisements on candidate evaluations (Brader, 2005; Carraro, 2010; Carraro and Castelli, 2010).

Chapter 4

The Impact of Implicit Racial Bias on Political Information Processing

4.1 Introduction

The consensual view in cognitive psychology and neuroscience is that information processing and attitude formation operate via a dual process model of thinking and reasoning (De Houwer and Moors, 2005; Evans, 2003; Evans, 2010). Specifically, a distinction is made between unconscious or implicit "system 1" and conscious or explicit "system 2" level information processing. System 1 is characterized as spontaneous, fast, and effortless processing, which operates below a conscious level of awareness. System 2, on the other hand, is described as processing that is deliberative, effortful and self-aware.

That spontaneously activated attitudes, relatively inaccessible to conscious thought, may potentially influence subsequent conscious deliberation and decision-making has garnered significant attention across the social sciences. Research has primarily focused on the potential impact of implicit attitudes on social behavior (Ferguson and Porter, 2010; Hassin et al., 2010; Peugini et al., 2010). Perhaps the most frequently studied implicit attitude (as it relates to social behavior) pertains to implicit racial bias. A plethora of studies have found implicit racial attitudes to be associated with a wide variety of behaviors including social interaction (Dovidio et al., 1997; McConnell and Leibold, 2001; Rudman and Ashmore, 2007), hiring decisions (Bertrand e al., 2005; Rooth, 2010; Ziegert and Hanges, 2005), medical diagnoses, (Green et al., 2007) and police behavior (Glaser and Knowles, 2008).

In political science, there is also a growing literature on the potential effects of implicit racial bias on political decision-making. Researchers exploring the determinants of anti-immigration attitudes in the U.S. have found implicit racial prejudice to be an important predictor of immigration policy preferences even when controlling for a number of known confounds (Perez, 2010; Malhotra et al., 2013; Knoll, 2013). In the study of candidate evaluation and vote choice, numerous studies have found implicit racial attitudes to be a significant predictor of voter decision-making. Researchers (Nosek et al., 2009; Payne et al., 2010; Pasek et al., 2010; Finn and Glaser, 2010; Mo, 2011) have shown implicit racial attitudes to be associated with lower favorability ratings and a decreased likelihood of voting for Barack Obama in the 2008 U.S. Presidential election, even when controlling for explicit forms of prejudice.

Despite a consensus beginning to emerge that implicit attitudes may be an important correlate for a variety of political outcomes, far less attention has been paid to examining the pathways through which implicit attitudes may influence political behavior. Specifically, it remains largely unclear the mechanism linking implicit attitudes to acts of political decision-making such as candidate evaluation or vote choice. Hence, the goal of this is paper is to systematically investigate one such potential mechanism to account for this association: **selective information search**. Building off of the work from Lodge, Taber, and colleagues (2000, 2005, 2006, 2013) on hot cognition, this paper directly examines the extent to which implicit attitudes bias information search in a direction consistent with one's implicit priors. Specifically, I explore whether implicit racial bias drives selective information search (and in turn, biased-decision-making) in the context of a political campaign that includes an African American candidate.

To address this hypothesis, I leverage the dynamic information board methodology (Lau and Redlawsk, 1997; 2006) to construct a mock presidential primary campaign where subjects are presented a matrix of "learnable" information about three fictional political candidates (including one African American candidate), which continuously flow down their computer screen for a limited length of time. At the end of the

campaign, subjects are then asked for their vote choice followed by a series of candidate evaluation items. Following the primary campaign, subjects are administered the "black-white" Implicit Association Test (IAT) as a measure of implicit racial bias.

By setting up the experiment in this way, this design provides a direct assessment of how implicit racial bias impacts political information search and voter decision-making. Specifically, I test whether implicit racial bias helps predict the comparability of search patterns across political candidates in terms of both the amount and duration of information accessed. Moreover, this design is also capable of examining the extent to which the association between implicit attitudes and political decision-making is mediated by selective information search.

In terms of results, I find that implicit racial bias has a consistent negative effect on information search for the African American political candidate. Specifically, I report that implicit racial bias has a negative effect on both the amount of opened information and reading time duration for the African American candidate. I report no such effect for the two white candidates in the primary campaign. Turning to candidate evaluation and vote choice, similar to other studies I show implicit racial bias to be negatively associated with both feeling thermometer scores and vote choice for the African American candidate. However, interestingly I find that this negative relationship loses significance once I account for information search patterns. Such findings suggest that one pathway through which implicit racial bias impacts voter decision-making is via biased information search.

These findings have a number of important implications for the growing literature on the relationship between implicit attitudes and political behavior. First, the findings from this paper provide systematic evidence consistent with Lodge and Taber's hot cognition hypothesis that implicit bias is an important source of motivated reasoning. Relatedly, by demonstrating the robust association between implicit attitudes and information search, this study sheds further light on how implicit attitudes indirectly impact political decision-making through information processing.

Finally, in recent years there have been a number of mixed findings pertaining to the effect of implicit racial bias once we account for explicit forms of racial prejudice. Some studies report a significant effect for implicit bias even when controlling for explicit racial attitudes (Perez, 2010; Nosek et al., 2009; Payne et al., 2009) while others fail to recover a significant effect (Ditonto, Lau, Sears, 2013; Ryan and Kinder, 2013; Segura and Valenzuela, 2010). One conclusion that some may draw from these recent mixed findings is that implicit attitudes provide little added explanatory leverage in certain domains of political behavior. However, the results from this study may help reconcile these contradictory findings. Specifically, previous studies may be underestimating the impact of implicit attitudes since they are assessing such effects at the end of the decision-stream (i.e. choice behavior). By highlighting the relationship between implicit bias and political information processing, my findings are consistent with the argument that implicit attitudes are likely to have a greater impact on spontaneous, non-verbal forms of behavior (Dovidio et al., 1997; Kinder and Ryan, 2013).

4.2 Literature Review

4.2.1 What are implicit attitudes?

Before briefly reviewing some of the studies that have explored the link between implicit attitudes and political behavior, let me first further define what I mean by an "implicit attitude" or system 1 processing. Social psychologists typically characterize an attitude as a memory-based affective association between a category (or object) and its evaluation (Fazio et al. 1982, 1986; Bargh, 1997). With so-called "implicit" attitudes, the affective association between an object and the evaluation of that object is activated in a relatively automatic fashion (Bargh, 1997; Fazio et al., 2007 Greenwald et al., 1998). By automatic, we generally mean that these affective evaluations are activated spontaneously, outside of one's full control or awareness, and expend few cognitive resources. This is in contrast to explicit attitudes or system 2 processing, which are considered cognitively effortful, demand attention, and are presumed to be based on intentional search of the working memory (Lodge and Taber, 2013).

While this distinction between system 1 and system 2 level processing or "implicit" and "explicit" attitudes is a simplification and of course there is quite a bit of overlap in use between the two systems; there is a fairly extensive literature suggesting that system 1 processing or implicit attitudes enter the decision-stream before any cognitive considerations come consciously to mind (Zajonc, 1980; 2000). Empirical evidence shows that object- evaluation associations are formed within 200-300 milliseconds upon exposure to an object, which is well before cognitive considerations are activated (Burdein, Lodge, and Taber, 2006; Le Doux, 1994, 1996; Morris et al., 2003). This has led some to characterize such implicit processing as an affect heuristic (Slovic, 1999; Slovic et al., 2004, 2007; Sniderman et al., 1991), which represents a quick and dirty pathway for the brain to simplify thought and decision calculations that precedes conscious deliberation.

4.2.2   Implicit Attitude Research in Political Information Processing

As mentioned earlier, researchers have amassed considerable evidence to suggest that implicit attitudes are potentially important predictors of a variety of political outcomes. Despite such evidence, it remains unclear the mechanism linking implicit attitudes to political decision-making. As Lodge and Taber (2013) point out, there are two ways that implicit processing can drive political behavior: (1) they may directly trigger a snap judgment or "gut reaction" that is entirely out of conscious awareness or thought (e.g. Todorov, 2005, 2007, 2010); (2) they may indirectly drive behavior through their impact on the process guiding political decision-making.

Todorov and colleagues (2005, 2007, 2010) have provided much of the research on the direct effect of automatic evaluations on political decision-making. In their studies, subjects are presented with pictures (facial images) of the winner and runner-up in congressional races and were asked to make a variety of trait judgments, including competence. The authors report that even with extremely short exposure times (100ms) for the facial images, subjects competency ratings correctly predicted election outcomes

at better than chance levels. Such findings provide strong evidence to suggest that pre-conscious judgements can be a strong predictor of behavior. However, as Verhulst et al (2010) point out, these pre-conscious judgements operate almost entirely through competency ratings thus suggesting that such effects can be mediated as subjects learn more information about candidates' credentials.

This has led many scholars to focus on some of the indirect pathways through which implicit attitudes may impact political behavior. Much of this work stems from Lodge, Taber, and colleagues' (2000, 2004, 2005) "hot cognition" research program, which contends that all sociopolitical concepts (e.g. political leaders, groups, issues, etc.) are affect-laden (positively or negatively) and that this affect is linked directly to memory, where it is automatically activated upon exposure to a concept. Subsequent research has focused on some of the downstream effects of the hot cognition hypothesis and how such automatic attitudes may bias deliberative behavior.

For instance Taber and Lodge (2006) find evidence consistent with the theory that affect-driven processing produces selective information exposure. In their experimental design, subjects were given an information board and asked to search for information about two sets of political issues. Subjects were also asked to assess the strength of eight arguments (four pro and four con) for each policy issue. Taber and Lodge report that subjects were far more likely to access information and support arguments congruent with their pre- existing beliefs (as measured by an attitude battery administered prior to the experiment). However, an important limitation to this study is that the authors do no directly measure subjects automatic affective associations and instead rely solely on self-reported policy preferences as a proxy for automatic affective associations.

In a more recent paper, Erisen, Lodge, and Taber (2013) address this critique and use the subliminal priming procedure to assess whether affective primes facilitate and/or inhibit subsequent conscious thought and deliberation. The authors find that subconsciously priming subjects with a positive or negative affective cue significantly altered

their retrieval of information in the LTM and how they consciously rationalized argument for or against a particular policy issue. Specifically, the priming condition significantly impacted the possibility of subjects citing (either negative or positive) reasons to support their argument.

Finally and most directly related to this paper, recent work by Galdi and colleagues (2012) attempt to investigate whether automatic associations trigger selective information search among those undecided about a particular policy. To test this, the authors administered a single category IAT that pairs positive and negative words with picture symbolizing the integration of Turkey into the EU. One week later subjects were given a selective exposure task where they were presented with pairs of newspaper headlines that suggested either a favorable or unfavorable article relating to the inclusion of Turkey into the EU. In line with their predictions, the authors find that, among undecided voters, selective exposure was significantly predicted by subjects' IAT scores. The authors conclude that their findings point to information search as a potential mediator for the link between implicit attitudes and individual decision-making.

4.2.3   Outstanding issues

In sum, while recent work has made some headway on the mechanisms linking implicit attitudes to political decision-making, there is still much work to be done. Galdi et al (2012) offers the best attempt at addressing this question by trying to empirically verify the link between implicit attitudes and selective information exposure, however they use only a single crude measure of information search in their study. Moreover, the type of implicit attitude under examination may also be an important factor in determining how such attitudes enter (and impact) conscious thought.

Previous research has shown that implicit racial bias influences behavior in often subtle and more complex ways than other types of implicit attitudes (Dovidio et al 1997). An important factor here is that individuals may be more motivated to control their implicit racial biases given its socially sensitive nature (Fazio and Towles-Schwen, 1997; Mo, 2011). This also may lead to implicit racial biases being (perhaps more

than other types of implicit attitudes) more likely to impact non-verbal behaviors than explicitly deliberative forms of behavior (Dovidio et al., 1997). Returning to the Galdi et al study on implicit attitudes and selective information exposure, the focus is on implicit attitudes toward Turkish integration into the European Union. Since such attitudes are not socially sensitive, their impact on deliberative behavior may be inflated since such attitudes are not subject to potential self-censoring (Fazio and Towles-Schwen, 1997).

The experimental design used in this paper offers several advantages compared to previous studies. First, I provide a more careful analysis of selective information exposure by leveraging a variety of different ways to measure information search. By using the dynamic process tracing methodology, I can assess, among other things, the comparability of search across alternatives, the total amount of information accessed, the content of information accessed, and the total duration of time spent viewing information. Secondly, by using the IAT to measure implicit racial prejudice, I provide a "hard" test for exploring the link between implicit attitudes and selective information search since socially sensitive attitudes are often subject to self-censoring. Finally, by examining not only information search but also voter decision-making, I am in a position to explore not only the relationship between implicit attitudes and information search, but also whether such biased information processing in turn impacts candidate evaluation and vote choice.

4.3 Experimental Design

4.3.1 Participants

This study was conducted in the fall of 2013. All of the experimental sessions were administered in the same laboratory in a large research university. A total of 115 students (59 male and 56 female) were recruited to participate in this study. Approximately one week prior to the lab sessions, subjects were administered a pre-study questionnaire, which included demographic items and two measures of explicit racial prejudice. Our demographic items include gender, race, and major. For the explicit racial prejudice

measure, I use Kam and Kinder's (2008) ethnocentrism scale as well as Kinder and Sear's (1982) symbolic racism scale.[1] Finally, I also include a measure of subjects' political ideology and party identification.[2]

Upon arrival to the lab, subjects are given a short political knowledge quiz before starting the mock presidential primary campaign. The political knowledge quiz consists of eight items and covers a variety of topics including questions on the branches of government, identifying well-known political figures, and foreign policy.[3] Once subjects completed the mock primary campaign and voted in the election, they moved to the final part of the study, where I administered the black-white race IAT.

4.3.2 Procedures

I use the dynamic process tracing methodology to construct a mock U.S. presidential primary election. The dynamic process tracing methodology uses an information board to track all decisions made by subjects when learning about different alternatives (in this case different candidates). Information about each candidate is presented to subjects on a computer screen and subjects must use the mouse to make selections as to what topics they are interested in learning more about. The information board contains six information boxes (at any one time), which continuously move down the screen.[4] Each box contains a label with a very short description of the nature of the information contained in the box and what candidate the information is for (e.g. "Lou Bakers Stance on Welfare"). Moreover, next to each information box is a photo of the candidate that the information box headline references.

---

[1]I administered the demographic and explicit prejudice items on week prior to subjects' lab sessions so as to avoid the potentially biasing effect of administering the explicit and implicit prejudice items in the same session.

[2]The text for all of these items can be found in the appendix.

[3]See the appendix for a list of the questions.

[4]Specifically, the six information boxes scroll continuously down the computer screen, with the box at the bottom disappearing and the new information box appearing at the top of the stack. Hence, every three seconds, subjects see the following: 1) a new information box appears; 2) all previously displayed information boxes move down one step; 3) one information box drops out of the information scroll area.

The order in which information is presented varies randomly as the subject proceeds through experiment. By clicking on the information boxes, subjects can learn more details about issue positions, candidate personalities, polls, endorsements, and background information about the candidates. When the subject finishes reading the information, they can return to the main screen where the information scroll continues to flow.[5] Figure 4.1 below provides a screenshot of how information is displayed to subjects on the computer screen.

- Enter Figure 4.1 Here -

The Candidates

In my mock U.S. presidential primary election subjects must choose to participate in either the Democratic or Republican primary. The choice of using a primary election serves to mitigate the overwhelming impact of party identification on information search and vote choice. In each primary, there are three candidates. The racial composition of the candidates is two white to one African-American.[6] In each primary, there are two moderate candidates and one highly partisan candidate. I assign policy positions to our candidates via independent coders ideology ratings of a wide range of policy stances. The composite average (i.e. the average of all issue position ratings for each candidate) was used to create an overall ideology score on a 1 to 7 scale. Below are the set of ideology scores for the three candidates in each primary election.[7]

---

[5]Please note that the information boxes continue to flow down the screen even when subjects click to open a particular information box. This is central to the idea of a dynamic (as opposed to static) information board, which forces subjects to make tradeoffs as to what information to look at.

[6]The decision to use three candidates stems from two concerns. First, the purpose of this study is to examine the impact of implicit racial bias (against African Americans). As such, I will need at least one African American candidate in our mock election. The issue is that if I only use two candidates (including one African American candidate), subjects may infer (at least partially) the nature of the experiment (i.e. that the study is in someway related to race). By using three candidates (including one African American candidate) for each primary election, the likelihood that subjects will make such an inference is much lower. Secondly, given that political campaigns tend to overwhelm voters with information, I wanted to do our best to replicate such an information environment while at the same time not making information so overly abundant that individuals are unable to retain any of the information that they were exposed to.

[7]Please note that these ideology scores were adapted from Lau, Redlawsk, and Andersen, (2008).

Table 4.1: Candidate Ideology Scores for Both Primaries

Candidate	Ideology Avg.	Off Median
Liberal Candidate	1.94	-2.06
Moderate Candidate A (Democrat)	3.77	-.23
Moderate Candidate B (Democrat)	3.41	-.59
Moderate Candidate A (Republican)	4.59	.59
Moderate Candidate B (Republican)	4.65	.65
Conservative Candidate	6.09	2.09

By design the two moderate candidates (in each primary election) differ in race (i.e. one white and one African American candidate). The highly partisan candidate is always white. To control for subtle differences between the two moderate candidates (in each primary), I randomly assign whether the African American candidate is "Moderate version A" or "Moderate version B." This washes out any potential differences in information search patterns and vote choice produced by the subtle differences in policy stances or personal characteristics between the two moderate candidates in each primary. In addition to blocking out differences in candidate ideology and personal characteristics, I also make some effort to control for the physical attributes of the candidates. All candidates are males, are of similar age (50 to 60), and are of similar physical attractiveness.[8]

The Information Environment

As for the structure of the information environment, each candidate has forty unique information attributes that appear (twice) during the campaign. These forty information attributes include twenty issue stands, party affiliation, political ideology, seven items referencing personal background information, six personality descriptions, a picture, two performance evaluations, and two items pertaining to their campaign strategy. The campaign also includes endorsements from nine interest groups and the results from

---

[8]All potential candidate images were rated by independent coders for attractiveness. Only images with similar attractiveness scores were assigned to candidates.

two different polls conducted at the early and late stages of the campaign. Candidate endorsements are randomized for the two moderate candidates and the poll results always show each candidate to be neck-and-neck. In terms of timing, a new information box enters the screen every three seconds and the campaign in total lasts for approximately sixteen minutes.

Immediately after the campaign, subjects are asked to cast their ballot for the primary nomination. Images of each candidate are presented side-by-side to remind subjects of each candidate. After the vote choice item, subjects are administered feeling thermometers for each candidate. Lastly, upon completion of the candidate evaluations, subjects are given a memory recall task where they are asked to write down as much information as they can remember about each candidate. Once this is finished, subjects move onto the IAT portion of the experiment.

4.3.3   The "Black-White" Implicit Association Test

Perhaps the best known and most widely used procedure to measure implicit attitudes is the IAT, or Implicit Association Test (Greenwald et al., 1998). Numerous studies (Bar-Anan and Nosek 2012; Nosek, Greenwald and Banaji 2007; Egloff, Schwerdtfeger and Schmukle 2005; Schmukle and Egloff 2004, 2005) have shown that IAT scores are arguably the most reliable. The IAT is a computer-based measure involving the rapid sorting of stimuli (words and facial images) into two key paired categories. Subjects complete many trials, where each trial requires the sorting of an individual stimulus item (usually words or faces that denote either "white" or "black"). Stimuli appear randomly and individually on the center of the computer screen and, as each stimuli appears, subjects must press the "E" or "I" key on a computer keyboard to quickly (and correctly) sort a given stimuli into the correct pairing ("white or good" and "black or bad" in some blocks of trials; and "white or bad" and "black or good" in other blocks of trials).

The logic behind the IAT is that individuals react more quickly to pairs of closely associated categories (Bravo and Greco, 2013). In contrast, individuals are slower to

react to pairs of categories that are not closely associated. This difference in reaction times (or the "IAT effect") is considered an indicator of the association between the categories used. If one category contains evaluative attributes (e.g. words characterized as positive or negative) and the other category contains target stimuli (e.g. African Americans and Whites, or men and women, etc), then the difference score in reaction time is seen as an indicator of an individual's attitude toward that target category.

The black/white IAT consists of seven blocks (see Table 3.1). The key blocks of interest to compute IAT scores are the blocks with "double-categorization" tasks, that is, blocks that contain paired categories: White is paired with either Good ("compatible") or Bad ("incompatible"); and Black is paired with Bad ("compatible") or Good ("incompatible"). Subjects are thus presented with a "compatible" sorting task, where: if the stimulus that appears on the computer screen is either a white facial image or a positive word, the subject must quickly press the "I" key; and if the stimulus is either a black facial image or a negative word, the subject must quickly press the "E" key). Subjects response times to these compatible sorting tasks are recorded. Subjects are then presented with the "incompatible" sorting task, where they perform the same operation, but now using flipped pairings ("White" is now paired with "Bad", and "Black" is paired with "Good"). Once again, subjects response times are recorded.[9]

To calculate an IAT score, I subtract the average response time of the "compatible" task from the average response time of the "incompatible" task to calculate the main quantity of interest: the "IAT effect". If subjects display slower average sorting times (and/or more mistakes) for the "incompatible" task than for the "compatible" task, then they are characterized as possessing a negative stereotype toward African Americans.[10]

---

[9]The order of sorting tasks (i.e. "compatible" or "incompatible") is randomized in the administration of the IAT; that is, some subjects will go through the "compatible" tasks first, whereas others will instead complete the "incompatible" tasks first.

[10]This not to say that an IAT score will capture all relevant dimensions of prejudice or stereotypes. Arkes and Tetlock (2004) have noted that IAT scores need not reflect personal animosity, but instead may well also reflect cultural knowledge about social norms/groups, although there is some disagreement on this point (Nosek and Hansen (2008) and Perez (2010).

A quick note on incentivizing the IAT

In this study, I use the novel approach of providing monetary incentives to subjects when administering the IAT. The motivation behind this design strategy is extensively detailed in Bravo and Greco (2013) and is also discussed in Chapter 3. In line with Bravo and Greco (2013), I provide a performance- based monetary incentive to IAT test-takers. Borrowing from Bravo and Greco (2013), in my experimental design all subjects are administered exactly the same standard race (black/white) IAT. I make no modifications to the actual IAT. However, unlike the standard IAT, in my studies all subjects are informed that they are eligible to earn extra money, depending on how quickly and accurately they complete their IAT. As discussed by Bravo and Greco (2013) this procedure is consistent with how the standard IAT is presented to test-takers, where subjects are actually asked to go as fast and accurately as possible when completing the sorting tasks that comprise the IAT.[11] Hence, all of the subjects in my experiment are told to go as "fast and accurately as they can when completing the IAT sorting tasks; however, in addition to these standard instructions, I also offer subjects a performance-based monetary incentive for doing so. Similar to Bravo and Greco (2013), subjects are informed that those in the top 20% (based upon both the speed and accuracy with which they complete the sorting tasks), will receive a $10 reward.

4.4   Results

For my analysis, I am in a position to examine both candidate information search and candidate evaluation as outcome variables. To examine the content of an individual's information search, I focus on two sets of variables: the amount of opened information boxes for each candidate and the total time spent reading information about each candidate. To account for differences across subjects in total amount of opened information and total reading time, I compute the variables as percentages. In other words, for each subject I compute the percentage of opened items related to each candidate.

---

[11]See the Appendix in Chapter 3 for exact wording of the instructions given to subjects

For example, the percentage of opened items related to candidate X for a given subject is computed as:

Total number of opened items for candidate X / Total number of opened items in the primary

This same calculation is done for the time spent reading information about each candidate. Again, this variable is created by simply taking the ratio between the total time spend reading information about candidate X and dividing the total time spent reading information about all of the candidates.[12] [13]

Another important thing to note pertains to how I will display the results. Since in both the Democratic and Republican primary, I use exactly the same setup (i.e. three candidates: one white moderate, one African American moderate, and one white partisan), I will aggregate the data across primaries so that I have three "meta" candidates:

(1) Candidate 1: Moderate, African American
(2) Candidate 2: Moderate, white
(3) Candidate 3: Partisan, white.

For simplicity and to maximize sample size, all results will be aggregated up to this level.[14] To account for potential differences in information search patterns between Democrats and Republicans, a partisanship control variable is included in all multivariate analysis.

---

[12]The rationale for computing a percentage as opposed to simply using the raw values is to account for variance in depth of information search that is unrelated to the experimental design. By calculating a percentage, I am constructing a measure of information search unique to each subject.

[13]It is also important to note that for both the percentage of opened items and reading time variables, all information items but for candidate endorsements is attributable to the one of the three candidates. Recall from chapter 3 that endorsements to the moderate candidates are randomly assigned. I report all results including the endorsement information items even though they are not candidate specific. However, as a specification check, all analyses were re-run excluding the endorsement items and there are no substantive changes to the results. Additionally, endorsement information items were randomized between the candidates, so this should have no effect on the information search results.

[14]Separate analyses were conducted at a more granular level (i.e. examining differences across candidates for Democrats and Republicans respectively), however this yielded no substantive differences in results.

4.4.1   Descriptive analyses

I begin by comparing the percentage of opened items for each candidate. Table 4.2 provides the mean and standard deviation scores for each candidate.

- Insert Table 4.2 Here -

The results here show that relative to the other candidates, the African American candidate maintained the lowest percentage of opened items. Moreover, when comparing the African American candidate to the white moderate candidate, I see a statistically significant difference at the .05 level. This result is noteworthy given that these two candidates are in expectation equivalent.[15] While there are some sizable differences in mean scores, the story is a bit different when I simply look at the distribution of scores for the moderate white and black candidates. Figure 2 reports the kernel density plot for the percentage of opened items for the two moderate candidates.[16]

- Insert Figure 4.2 Here -

As figure 4.2 shows, there is actually little difference in the distribution functions across the two candidates. A Kolmogorov-Smirnov test for the equality of distributions fails to be rejected at the .10 level. In sum, in my initial comparison of the percentage of opened items across candidates, I report mixed findings as to the differences between the white and African American moderate candidates. While I find some differences in the mean scores, the overall distribution of scores do not look much different across candidates.

However, what happens once I account for differences in implicit racial bias within the sample? A simple way to first examine the impact of implicit racial bias on the

---

[15]Recall that the moderate candidate versions were randomized.

[16]I report density plots to examine whether there are differences in the distribution of opened items between the two candidates.

percentage of opened items across candidates is to split the sample into two groups: "high" vs. "low" levels of implicit prejudice. A crude way to do this is to use a median split. Malhotra et al (2013) use a similar strategy in their descriptive analysis of the IAT effect. In other words, I can characterize those with IAT scores above the median as having relatively high levels of implicit bias while those below the median having relatively low levels of such bias. Figure 4.3 shows a kernel density plot for the percentage of opened items for the African American candidate now accounting for "low" and "high" IAT subjects.

- Insert Figure 4.3 Here -

Here I see sizable differences in the percentage of opened items for the African American candidate with high IAT subjects accessing significantly fewer items than low IAT subjects. A Kolmogorov-Smirnov test for the equality of distributions is rejected at the .05 level. The differences between high IAT and low IAT subjects is further evidenced by figure 4.4, which presents the mean score (and its associated confidence interval) for percentage of opened items for the African American candidate.

- Insert Figure 4.4 Here -

Again, I see statistically significant differences in the point estimates for the two IAT groups. High IAT subjects opened over 3% less information items for the African American candidate. This is statistically significant at the .10 level. As a robustness check, I can also examine whether I see differences between the low IAT and high IAT groups when examining the percentage of opened items for the white moderate candidate. In other words, it may be the case that the differences between the low and high IAT groups reflects something unrelated to racial bias (maybe simply differences in processing speed?). Figures 4.5 and 4.6 present the kernel density and confidence interval plots for the white moderate candidate.

- Insert Figures 4.5 and 4.6 Here -

Here I see that there are no statistically significant differences between the low and high IAT groups for the white moderate candidate. As expected, the differences between the low and high IAT group only appears to operate for the African American candidate.

Let us now shift attention to the second information search variable, which measures the percentage of reading time spent on each candidate. Table 4.3 presents the summary statistics for this measure.

- Insert Table 4.3 Here -

I see a very similar pattern to that of the opened item variable. Subjects spent the least amount of time reading about the African American candidate and again the difference between the African American and white moderate candidates is statistically significant at the .05 level. Figure 4.7 examines the density plot of the reading time variable for the two moderate candidates and again I see that there are some differences, however these fail to reach statistical significance.

- Insert Figure 4.7 Here -

Turning to the IAT median split comparison, I can again examine whether large differences emerge once I account for implicit racial bias. In examining these plots, I see that subjects in the high IAT group spent significantly less time reading about information pertaining to the African American candidate. Comparing the African American candidate to it's white moderate counterpart, I see statistically significant differences in both distribution functions as well as the mean percentage of reading time for each candidate.[17]

---

[17]Again, as a robustness check I can examine both density distributions and confidence interval plots for percentage of reading time for the white moderate candidate. While I do not show these results in this chapter, I find that there are no differences between the low and high IAT groups for the white moderate candidate.

- Insert Figures 4.8 and 4.9 Here -

To sum up, at first glance I find evidence consistent with the hypothesis that implicit racial bias is associated with selective information search.[18] Specifically, I see that subjects with relatively high levels of implicit racial prejudice access less information pertaining to the African American candidate and spend less time reading such information when opened. While this initial evidence is persuasive, the use of a median split variable is somewhat arbitrary and does not account for potential confounds. Thus, a more careful multivariate analysis is needed to examine whether implicit racial bias impacts selective information search when controlling for several known covariates of information search.

4.4.2 Regression analysis

Let's now turn to modeling candidate information search as a function of implicit prejudice as well as some additional covariates.[19] As mentioned earlier, subjects were administered a list of survey items that included demographics, explicit racial prejudice, political ideology, and political knowledge. The key question is whether I can still recover an IAT effect once I include a number of covariates that may also be associated with information search patterns. Specifically, I will report results from two regression models. In the first model, I regress the total percentage of opened items for the African American candidate on IAT scores and a number of covariates. In the second model, I use the same model specification, however the outcome variable this time is total percentage of reading time for the African American candidate.

---

[18]Additionally, I tested whether I find similar differences in the information search items for explicit racial bias. Specifically, I use the symbolic racism scale to create a median split variable for explicit racial bias. The results for this analysis can be found in the appendix figures 1 and 2. While, I find results consistent with what is reported for implicit racial bias, these results failed to reach statistical significance.

[19]For all analyses, I use robust standard errors. Alternative models without robust SEs yield no substantive difference. Results from these models are available upon request.

Predictor variable

For all of the models, I enter IAT scores as a continuous predictor. As mentioned earlier, IAT scores are computed by taking the difference between the incompatible and compatible blocks, with positive difference scores indicating higher levels of implicit bias. Hence in my models, IAT scores are simply raw difference scores that are log transformed.[20] The reason for using a log transformation is documented by Greenwald et al (2003) and centers on the strong right skew typically found in IAT data. By log transforming the IAT data, I reduce the effect of potentially large positive outliers, which may artificially produce large IAT effects.[21]

Controls

Turning to my covariates, to account for potential demographic differences, I include dummy variables for both gender and race (i.e. white vs. non-white). A second potential covariate may relate to differences in political ideology. Remember, since I use political primaries, I am able to control for the effect of partisanship within the mock election campaigns. However, since I am aggregating the data across primaries for the analysis, I need to account for potential differences in political ideology that may lead to systematic differences in information search patterns.[22] To do this, I collapse the political ideology survey item into three categories: liberal, conservative, and moderate. In my models, I use moderate as the baseline category.[23] Finally,

---

[20]It is important to note that I do not use IAT-D scores in my analysis. As detailed by Bravo and Greco (2013), the monetary incentive procedure outperforms the IAT-D in both internal and temporal reliability. More importantly, as Bravo and Greco (2013) note, using the IAT-D measure in conjunction with the monetary incentive procedure produces biased results since the monetary incentive (by design) reduces the block standard deviations, which serves as the denominator in the IAT-D scoring algorithm. As a result, by reducing measurement error, the monetary incentive artificially inflates IAT-D scores.

[21]Please note that a data appendix is available, which re-runs all of the main analysis presented in this paper with alternative IAT measures including the IAT-D and median block difference scores. There are no substantive difference in results using these alternative measures.

[22]In other words, the intercepts for the IAT effect may be different for liberals, moderates, and conservatives.

[23]I chose to use the political ideology scale in lieu of party identification simply because the political ideology scale maintained more variance across categories, which thus allowed for more meaningful comparisons in my models. However, all analyses were also run using a party identification categorical variable and the results show no substantive differences. These results are available upon request.

in addition to political ideology, I also control for differences in political knowledge. The reason for doing this stems from Lau and Redlawsk's (2006) finding that political sophisticates tend to engage in more shallow information search than non-sophisticates. To account for this, I include an additive political knowledge scale constructed from the political knowledge quiz administered prior to the primary campaign.

Finally, the most potentially important covariate of information search is explicit racial prejudice. Numerous studies report that both implicit and explicit measures of prejudice maintain significant effects on political behavior (Nosek et al., 2009; Payne et al., 2009; Finn and Glaser, 2010) while others find that explicit racial prejudice actually maintains much larger effects on behavior (Ditonto, Lau, and Sears, 2013; Kinder and Ryan, 2013). To control for the potential effect of explicit racial prejudice, I use Kam and Kinder's (2008) ethnocentrism scale to construct a difference score for items relating to whites and blacks. As Kam and Kinder (2008) note, an ethnocentrism score is produced by first creating an additive scale for all items related to the in-group and out-group respectively. Then a difference score is computed by subtracting the out-group value from the in-group value. Hence, for my explicit prejudice measure I compute a difference score where a positive score indicates higher levels of explicit prejudice against African Americans.[24]

4.4.3 Findings

Information search

Table 4.4 presents the results for the first model where I regress the percentage of opened items for the African American candidate on IAT scores and then sequentially include my control variables.

- Insert Table 4.4 Here -

---

[24]Please note that all of my models were also run using the symbolic racism scale and there are no substantive differences in the results. Please see the appendix for details. I chose to use the ethnocentrism scale because there is no missing data with this measure while for the symbolic racism scale, I needed to impute scores for twenty five missing values.

Here I see that the implicit racial bias measure maintains a statistically significant negative effect on the percentage of opened items for the African American candidate across all model specifications. Surprisingly, I report little effect for the explicit prejudice measure, which fails to reach statistical significance in any one of my model specifications. Interestingly, the race dummy also has a statistically significant negative impact on the percentage of opened items and this effect does not appear to run through implicit or explicit prejudice. For my final model specification, I include not only all of the covariates but also a measure of the total number of information items opened. I include this variable to control for potential differences across subjects with respect to depth of information search. In other words, some subjects may simply be more shallow information searchers than others, which may account for some of the variance in the percentage of opened items for the African American candidate. However, I find that even when controlling for the number of information items opened, I still recover a negative effect for implicit racial bias.

I can employ a very similar modeling strategy for the second information search variable, the percentage of time spent reading information pertaining to the African American candidate. Table 4.5 presents a series of models for this outcome variable.

- Insert Table 4.5 Here

These models differ slightly from my earlier models in that I include a measure of reading speed to account for differences across subjects in how quickly they read through the information boxes.[25] Similar to previous models, I also include a measure to account for the total amount of reading time spent on all of the candidates. Turning to the results, my findings are quite similar to the first set of models for the percentage of opened items. Again, I find that the implicit racial bias measure maintains a statistically negative effect across all model specifications while I fail to recover any sort of

---

[25]The reading speed variable was created by dividing the sum of the total word count for the three instruction screens prior to the mock campaign and the total time spent (in seconds) reading the three instruction screens.

meaningful effect for the explicit racial prejudice variable.

In sum, it appears that implicit racial attitudes have a consistent effect on candidate information search, even when controlling for a number of known confounds. Specifically, I find that implicit attitudes indeed contribute to selective information exposure with higher implicit prejudice folks searching for and reading less information about the African American candidate. In many ways, these findings are consistent with previous research highlighting that implicit attitudes are likely to have larger effects on non-verbal forms of behavior where subjects are less motivated to suppress such attitudes (Dovidio et al., 1997). Since this study focuses primarily on more subtle forms of behavior (i.e. comparability of information search across candidates), the effect of implicit bias is likely less vulnerable to self-censoring. This may also help to explain why I fail to recover any meaningful effect for the explicit prejudice measure.

That being said, one way to examine potential self-censoring or the "over-riding" of implicit attitudes is to examine whether I see an interaction between implicit bias and political sophistication. Several researchers contend that highly motivated individuals are able to override their implicit biases in order to make decisions. For instance, Mo (2011) shows that the effect of implicit racial bias on voter decision-making is strong for weak partisans but close to non-existent for strong partisans. Similarly, Kam (2007) shows that a cue of a candidate's party identification can help override the effect of implicit racial bias. However, other researchers contend the exact opposite and argue that implicit associations are more fully crystallized for political sophisticates thus increasing the likelihood that such individuals rely on their implicit attitudes (Lodge and Taber, 2005).

To address these competing claims, let's first regress the percentage of opened items for the African American candidate on IAT and all of the covariates and include an interaction between IAT and the political knowledge variable. To ease the interpretation of the interaction model, I use the dichotomized IAT median split variable although its important to note that there are no substantive differences in results when I use a

continuous IAT measure. Table 4.6 presents the results for this model.

- Insert Table 4.6 Here

Using this model specification, I am able to recover a significant interaction effect between IAT score and political knowledge. However, in order to understand the nature of this interaction I need to plot this effect. Figure 4.10 presents the conditional effect of IAT on the information search variable across one standard deviation above and below the median political knowledge score.

- Insert Figure 4.10 Here

When I plot the interaction effect, I see that implicit racial bias has a statistically significant effect on those below or at the median political knowledge score. However, when I move to those subjects above the median political knowledge score, I see that IAT has little effect. These findings are consistent with the hypothesis that political sophisticates rely less on their implicit biases to inform their decision-making.[26]

Candidate evaluation and vote choice

For the second half of our analysis, I can turn my attention to how implicit prejudice influences candidate evaluation and vote choice. At the end of mock campaign, subjects were first asked to vote for their most preferred candidate. After casting their ballot, subjects were then administered feeling thermometer (FT) items for each candidate and were given a memory recall task where they were asked to write as much as they can remember about each candidate. Focusing first on candidate evaluation, I can compute a feeling thermometer difference score by subtracting the FT score for our African American candidate from the FT score for the white moderate. To explore whether

---

[26]I ran the exact same analysis for the other information search variable, total percent of reading time spent on the African American candidate, and found exactly the same interaction effect. The regression results along with the plot of the interaction plot can be found in the appendix.

implicit racial bias influences candidate evaluation, I can first use the IAT median split variable to examine FT difference scores.

Figure 4.11 presents the kernel density plots for the FT difference score across the IAT median split variable. Here I see sizable differences in the distribution functions, with high IAT subjects rating the white moderate candidate higher than the African American candidate.[27] The null hypothesis of equality of distributions is rejected at the .05 level. Moreover, a Wilcoxon signed rank test shows a statistically significant difference at the .05 level.

- Insert Figure 4.11 Here

Again, I can examine the robustness of these findings by employing some multivariate analyses. Specifically, I can regress my FT difference score on the continuous measure of IAT and again sequentially include the covariates. However, in addition to including the typical covariates, I can also include some of the information processing variables in order to examine whether IAT's effect on candidate evaluation appears to be channeled through information search. Table 4.7 presents a series of models for the FT difference score.

- Insert Table 4.7 Here

In model 1, I present my baseline regression model that includes IAT along with the standard set of covariates. Similar to previous studies, I also report a negative association between implicit prejudice and feeling thermometer rating for the African American candidate. Somewhat surprisingly, I again are unable to recover a significant negative effect for the explicit racial prejudice measure, however the coefficient is in the expected direction. Once again I find that the race dummy is a significant predictor

---

[27]A larger FT score indicates greater favorability toward the African American candidate relative to the white candidate.

of subject's voter behavior with white subjects rating the white moderate candidate far higher then its African American counterpart, however this does not appear to be channeled through the prejudice measures.

In model 2, I examine whether I can still recover IAT's effect on the feeling thermometer difference score once I control for how much information subjects are able to remember about each candidate. Here, I introduce two new variables that account for the number of information items recalled for both the African American candidate and the white moderate candidate during the memory recall task.[28] The mean number of unique information recalled for the African American candidate was 3.99 while for the white moderate candidate it was 3.77. This difference was not statistically significant. As expected, I find that the two memory recall variables have the opposite effect on the feeling thermometer difference score. Specifically, the more items recalled about the African American (white) candidate has a positive association on candidate evaluation. However, more surprisingly I see that IAT score maintains a negative effect on candidate evaluation (that is very close to reaching statistical significance) even when controlling for memory recall.

Such a findings is consistent with a number of studies emphasizing the primacy of affect in decision-making. To better understand this, I can explore possible interaction effects between implicit attitudes and memory recall for the African American candidate. To the extent that crystallized implicit attitudes precede (and largely influence) deliberative thought, I should see memory recall having a larger effect on candidate evaluation for those subjects with lower levels of implicit bias. For those with higher levels of implicit bias, I should see the opposite effect with memory recall having little effect. The reason for this being that the primacy of affect for high implicit bias individuals is so strong that memory recall either has no overriding effect or such memory

---

[28]The criteria used for coding the opened-ended memory recall items for each candidate is as follows. Each unique piece of candidate-specific information was coded as one memory recall item. The accuracy of the memory recall items was not considered.

activation is biased in the direction of the implicit prejudice. Table 4.8 shows the results for the candidate evaluation regression, which now includes an interaction term for IAT and memory recall for the African American candidate. Moreover, figures 4.12 and 4.13 plot the interaction effect for low and high IAT subjects respectively.

- Insert Table 4.8 Here

- Insert Figures 4.12 and 4.13 Here

As expected, I find that for high levels of implicit bias, memory recall for the African American candidate has virtually no effect on candidate evaluation. However, for those with lower levels of implicit bias, I see memory recall having a (over-riding) positive effect. Such findings provide added evidence to those studies emphasizing the primacy of automatic attitudes (Lodge and Taber, 2013).

Finally, we can turn our attention to model 3 of Table 7, which includes not only the memory recall variables but also one set of the information search variables, total percentage of opened items for each candidate. As hypothesized earlier, if IAT's effect on voter decision-making operates primarily through candidate information search, then I expect that the direct effect of IAT should significantly weaken once I include a measure of information search into the candidate evaluation model. Examining my findings from model 3, I see that our information search variables have a statistically significant effect on candidate evaluation in the expected direction. More importantly, I see that coefficient for the IAT variable significantly decreases while it's accompanying standard error increases in size. This finding is consistent with the story that the effect of implicit racial bias on candidate evaluation is potentially mediated by selective information search.[29]

---

[29]Of course, I cannot demonstrate the validity of this conjecture empirically since a mediation analysis would need to be conducted. Moreover, as Imai et al (2012) note, to assess potential mediation, the proposed mediator must be exogenous, which in my case it is surely not true.

4.4.4 Voting behavior

For my final set of analyses, I can examine the effect of implicit racial bias on vote choice. Table 4.9 provides the descriptive statistics for vote choice.

- Insert Table 4.9 Here

Again, I see that the African American candidate fared the worst in the election. Moreover, the difference in vote share between the African American and moderate white candidate is statistically significant at the .05 level. Turning to the multivariate analysis, I can specify a model for vote choice by dichotomizing my vote measure. Specifically, in order to maximize sample size I construct a dichotomous measure of vote choice where I assign a value of "1" for a vote for my African American candidate and a "0" otherwise. Using this outcome variable, I can use the same modeling strategy (but instead of OLS, I use a logit specification) as with candidate evaluation and model vote choice as a function of IAT and covariates, then add in the memory recall variables, and then finally include the information search variables.

- Insert Table 4.10 Here

Table 4.10 presents the results from the three models of vote choice. Consistent with other studies, I find that IAT has a negative effect on the likelihood of voting for the African American candidate. However, once again I fail to recover an effect for the explicit prejudice variable (although it is in the expected direction). Similar to the candidate evaluation models, I again find that IAT's effect is robust to the inclusion of the memory recall items. However, like the previous models for candidate evaluation, I again see that IAT's effect dramatically weakens upon the inclusion of the information search variables, thus providing added evidence to suggest that the effect of implicit racial bias on vote choice is operating through biased information search.

4.5   Conclusion

Previous studies that have explored the determinants of selective information processing have either neglected to measure implicit attitudes directly or have used crude measures of information search as their outcome variable. Using both the Implicit Association Test (Greenwald et al., 1998) and the Dynamic Process Tracing methodology (Lau and Redlawsk, 1997, 2006), I constructed a mock Presidential primary campaign to test whether implicit racial attitudes were associated with systematic differences in both information search patterns and voting behavior across subjects. The results from this study show implicit racial attitudes are indeed linked to selective information processing. Controlling for a variety of confounds, I consistently find that implicit racial prejudice is associated with more shallow information search for an African American candidate. Specifically, implicit racial bias is negatively associated with both the percentage of opened items for the African American candidate as well as the percentage of time reading information about the African American candidate.

In addition to empirically assessing the link between implicit racial bias and candidate information search, I also was able to replicate previous findings regarding the effect of implicit prejudice on candidate evaluation and vote choice. However, unlike previous studies, in addition to demonstrating this relationship, this paper also sheds some light on the potential pathways through which implicit bias enters into voter decision-making. The results from my study show that the effect of implicit racial bias on candidate evaluation and vote choice significantly weakens once I account for candidate information search. Such findings provide support for the hypothesis that the effect of implicit attitudes on political behavior operates through information processing channels such as information search.

In sum, these findings have important implications for much of the work in political psychology and public opinion on the impact of implicit attitudes on voter decision-making. By providing important insight into the relationship between implicit attitudes and information processing, this study helps to unpack how automatic associations can

influence conscious beliefs and subsequent decision-making. Moreover, the results from this project help inform the growing debate on the significance of implicit attitudes as predictors of political behavior. While there has been recent push back by some scholars that implicit attitudes add little additional explanatory power to models of voting behavior (Ditonto, Lau and Sears, 2013; Kinder and Ryan, 2013; Segura and Valenzuela, 2010), such findings may be underestimating the influence of implicit attitudes by focusing only on direct effects. The findings from this study suggest that implicit attitudes impact more subtle forms of political behavior, specifically how individuals search for information about politics. Such findings are consistent with previous studies that have found that implicit attitudes tend to have stronger effects on nonverbal forms of behavior (Dovidio et al., 1997).

4.6   Tables and Figures

Trial Blocks in the Black/White IAT

No. of trials	Block	Items assigned to left-key response	Items assigned to right-key response
20	Demo 1	African American images	European American images
20	Demo 2	Unpleasant words	Pleasant words
20	Compatible Block 1	Unpleasant + African American images	Pleasant + European Americans images
40	Compatible Block 2	Unpleasant + African American images	Pleasant + European Americans images
20	Demo 3	European American images	African American images
20	Incompatible Block 1	Unpleasant + European American images	Pleasant + African American images
40	Incompatible Block 2	Unpleasant + European American images	Pleasant + African American images

Note: The order of sorting tasks (i.e. "compatible" or "incompatible") is randomized in the administration of the IAT. However, it is always the case that Compatible block 1 (Incompatible block 1) precedes Compatible block 2 (Incompatible block 2).

Table 4.2: Summary statistics for % opened items across candidate

Candidate	Mean	Std. Dev.
Moderate African American candidate	0.312[a]	0.064
Moderate White candidate	0.347[a]	0.069
Partisan White candidate	0.340	0.084
N		115

[a] A Wilcoxon signed rank test shows a statistically significant difference at the .05 level

Table 4.3: Summary statistics for % of total duration

Variable	Mean	Std. Dev.
Moderate African American candidate	0.309[a]	0.087
Moderate White candidate	0.349[a]	0.094
Partisan White candidate	0.341	0.114
N		115

[a] A Wilcoxon signed rank test shows a statistically significant difference at the .05 level

Table 4.4: Regression for % of information opened for African American candidate

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
IAT	-0.075**	-0.075**	-0.079**	-0.062**	-0.059*	-0.069**	-0.072**
	(0.032)	(0.032)	(0.034)	(0.031)	(0.031)	(0.033)	(0.034)
Explicit Bias		0.0002	0.0004	0.0003	0.00003	- 0.00003	0.001
		(0.001)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Female			-0.010	-0.008	-0.010	-0.011	-0.012
			(0.012)	(0.012)	(0.013)	(0.013)	(0.013)
White				-0.021*	-0.021*	-0.025**	-0.023*
				(0.011)	(0.011)	(0.012)	(0.012)
Political Knowledge					-0.001	-0.001	-0.001
					(0.002)	(0.002)	(0.002)
Conservative						0.030	0.030
						(0.022)	(0.022)
Liberal						0.016	0.017
						(0.018)	(0.018)
Total Opened							0.031*
							(0.018)
Constant	0.330***	0.329***	0.346***	0.353***	0.363***	0.354***	0.212***
	(0.009)	(0.009)	(0.024)	(0.024)	(0.315)	(0.029)	(0.08)
N	115	115	115	115	115	115	115
R^2	0.037	0.037	0.042	0.067	0.069	0.089	0.111

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]The dependent variable is calculated as a percentage. Specifically it is calculated by dividing the total amount of information opened for the African American candidate by the total amount of information opened for all candidates.
[b]The IAT variable is constructed by log transforming each trial score and then differencing B3, B4 from B6, B7.
[c]The Conservative and Liberal dummies were constructed from the standard ideology scale given to subjects. The baseline category is moderate.
[d]I use a log transformation for the total opened variable given it's strong right skew. All coefficients are unstandardized.

Table 4.5: Regression for % of reading time for African American candidate

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
IAT	-0.093**	-0.094**	-0.097**	-0.076✝	-0.074✝	-0.087*	-0.089*	-0.088*
	(0.047)	(0.047)	(0.049)	(0.049)	(0.049)	(0.049)	(0.049)	(0.049)
Explicit Bias		0.001	0.0008	0.001	0.001	0.001	0.001	0.001
		(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)
Female			-0.009	-0.007	-0.009	-0.010	-0.012	-0.012
			(0.017)	(0.017)	(0.018)	(0.018)	(0.018)	(0.019)
White				-0.025✝	-0.025✝	-0.030*	-0.030*	-0.028✝
				(0.017)	(0.017)	(0.018)	(0.018)	(0.018)
Political Knowledge					-0.001	-0.001	-0.001	-0.001
					(0.003)	(0.003)	(0.003)	(0.003)
Conservative						0.037	0.038	0.038
						(0.030)	(0.030)	(0.030)
Liberal						0.028	0.027	0.027
						(0.023)	(0.023)	(0.023)
Reading Speed							-0.010	-0.006
							(0.018)	(0.019)
Total Duration								0.033
								(0.029)
Constant	0.331***	0.330***	0.344***	0.353***	0.361***	0.345***	0.413***	0.180
	(0.013)	(0.013)	(0.033)	(0.032)	(0.043)	(0.043)	(0.123)	(0.250)
N	115	115	115	115	115	115	115	115
R^2	0.031	0.032	0.035	0.053	0.055	0.077	0.075	0.083

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]The reading speed variable was calculated by dividing the subjects' time spent on reading the three sets of directions prior to the experiment by the total number of words for the directions.
[b]I use a log transformation for the total duration and reading speed variables given their strong right skew.

Table 4.6: IAT's effect on % opened for African American candidate across political knowledge

Median IAT	-0.076**
	(0.027)
Political Knowledge	-0.006†
	(0.003)
Median IAT * Political Knowledge	0.009†
	(0.005)
Explicit Bias	0.001
	(0.003)
Female	-0.012
	(0.013)
White	-0.024*
	(0.012)
Liberal	0.019
	(0.018)
Conservative	0.033
	(0.022)
Total Opened	0.032†
	(0.018)
Constant	0.231**
	(0.083)
N	115
R^2	0.147

Table 4.7: Regression: Candidate Evaluation, Including Info Processing Variables

	(1)	(2)	(3)
IAT	-24.098†	-26.555†	-16.905
	(17.832)	(16.734)	(17.654)
Explicit Bias	-1.219	-1.065	-0.845
	(1.109)	(1.095)	(0.920)
Female	-1.124	-0.033	2.740
	(5.478)	(5.086)	(5.048)
White	-11.391**	-9.748*	-5.735
	(5.612)	(5.569)	(5.048)
Liberal	5.880	6.523	3.250
	(6.454)	(5.947)	(5.706)
Conservative	7.911	10.252	5.322
	(8.914)	(8.489)	(8.333)
Political Knowledge	-0.787	-0.753	-0.218
	(1.147)	(1.115)	(1.072)
Memory Count African American		3.793***	1.140
		(1.277)	(1.532)
Memory Count White Moderate		-4.373***	-1.826†
		(1.019)	(1.204)
% Opened African American			1.459***
			(0.406)
% Opened White Moderate			-0.696*
			(0.358)
Constant	15.369	13.524	-15.894
	(12.317)	(11.870)	(20.610)
N	115	114	114
R^2	0.091	0.212	0.320

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]The memory count variables are constructed from the memory recall tasks that subjects were given at the end of the experiment. Each unique item listed for a given candidate was coded as one memory. The accuracy of memories have not yet been coded.

[b]The % Opened African American variable is the percentage of opened items that were about the African American relative to the total amount opened.

[c]The coefficients for IAT in these models appear unusually large due to the fact that IAT is log transformed and the outcome variable is based on a 100 point scale.

Table 4.8: Interaction between IAT and Memory Recall on % FT Diff Score

Median IAT	-0.041
	(10.535)
Memory Count Moderate African American	2.194
	(1.709)
Median IAT * Memory Count African American Moderate	-2.853
	(2.511)
Controls	Yes
Constant	-1.978
	(12.316)
N	114
R^2	0.12

Table 4.9: Vote share across candidates

Candidate	Vote share
Moderate African American candidate	20.00[a]
Moderate White candidate	31.30[a]
Partisan White candidate	46.96
Wouldn't vote for any candidate	1.7
N	115

This difference is statistically significant at the .10 level.

Table 4.10: Logistic Regression: Vote for African American candidate

	(1)	(2)	(3)
IAT	-2.319✝	-3.521*	-3.046
	(1.621)	(1.901)	(2.535)
Explicit Bias	-0.095	-0.145	-0.320
	(0.147)	(0.189)	(0.255)
Female	0.248	0.297	0.239
	(0.529)	(0.615)	(0.754)
White	-0.389	-0.172	0.128
	(0.511)	(0.594)	(0.765)
Liberal	-0.911✝	-0.616	-0.872
	(0.579)	(0.700)	(0.971)
Conservative	-1.124	-0.872	-2.620**
	(0.811)	(0.963)	(1.378)
Political Knowledge	0.128	0.124	0.168
	(0.113)	(0.143)	(0.184)
Memory Count African American		0.797***	1.100***
		(0.197)	(0.346)
Memory Count White Moderate		-0.567***	-0.910***
		(0.173)	(1.204)
% Opened African American			0.280***
			(0.078)
% Opened White Moderate			0.147*
			(0.090)
Constant	-0.961	-2.654✝	-17.283***
	(1.290)	(1.675)	(4.753)
N	115	114	114
R^2	0.085	0.293	0.511

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]The outcome variable here is 1 if subject voted for the African American and 0 if they voted for the two other candidates or abstained.

Figure 4.1: Screenshot of Dynamic Process Tracing Environment

Figure 4.2: Density plot: % of opened items for the two moderate candidates
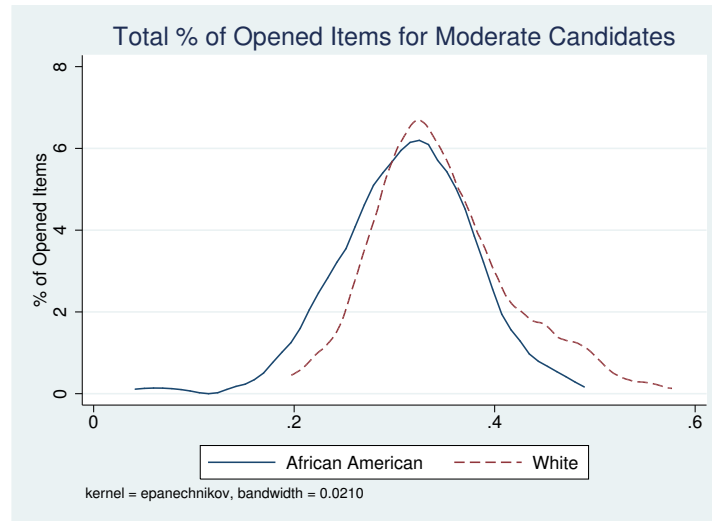
Figure 4.3: Density plot: % of opened items for African American candidate across IAT
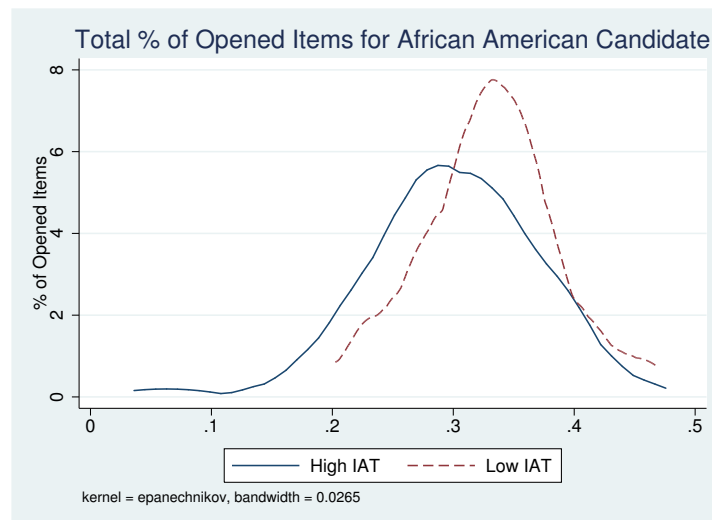
A Kolmogorov-Smirnov test for equality of distribution functions is rejected at the .01 level

Figure 4.4: CI plot: % of opened items for African American candidate across IAT
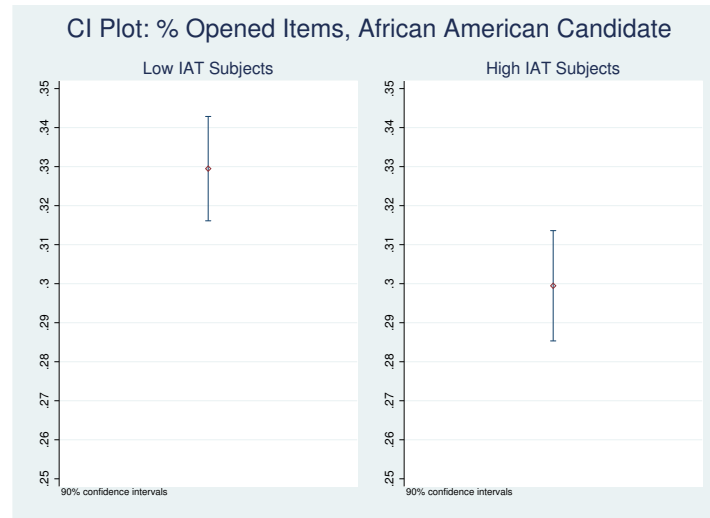
A Mann-Whitney difference of means test shows a statistically significant difference at the .05 level

Figure 4.5: Density plot: % of opened items for white moderate candidate across IAT
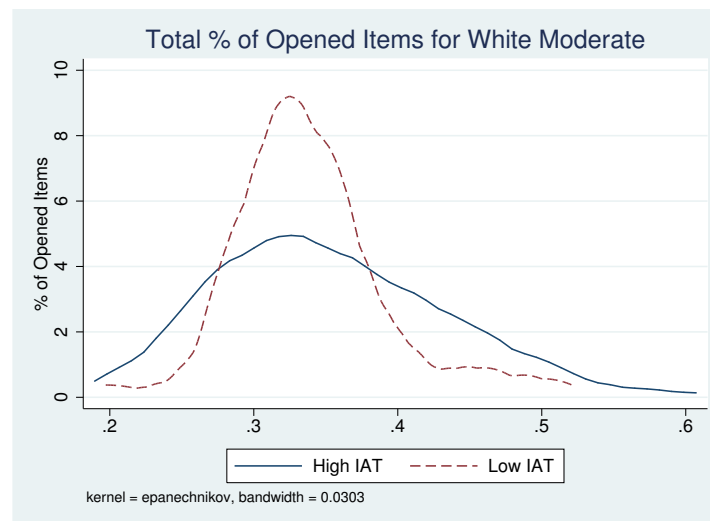
A Kolmogorov-Smirnov test for equality of distribution functions fails to be rejected

Figure 4.6: CI plot: % of opened items for white moderate candidate across IAT
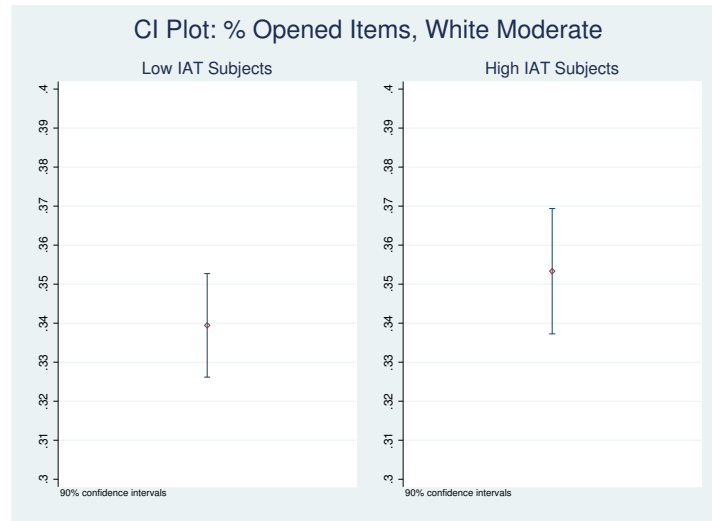
A Mann-Whitney difference of means test shows does not show a statistically significant difference

Figure 4.7: Density plot: % of reading time for the two moderate candidates
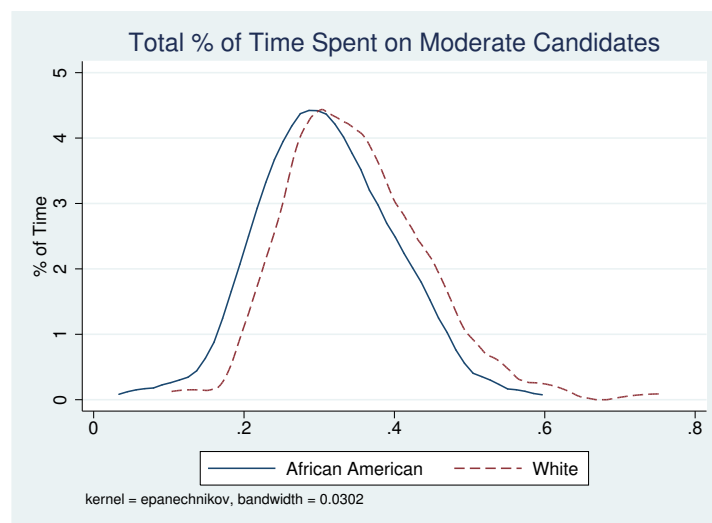
Figure 4.8: Density plot: % of reading time for African American candidate across IAT
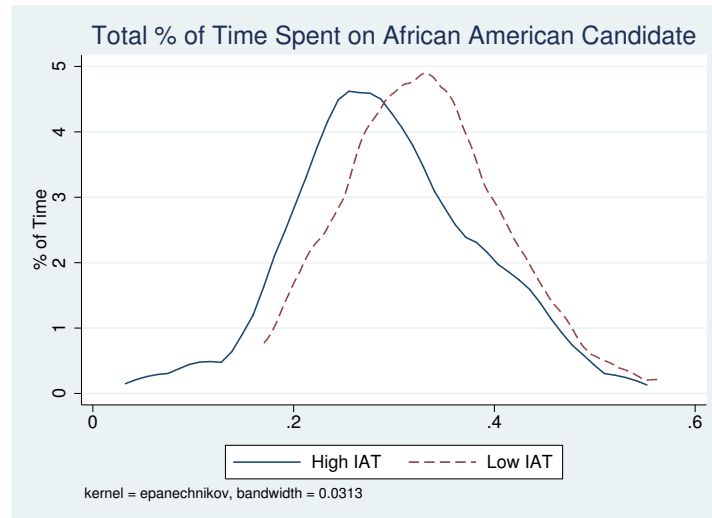
A Kolmogorov-Smirnov test for equality of distribution functions is rejected at the .05 level

Figure 4.9: CI plot: % of reading time for African American candidate across IAT
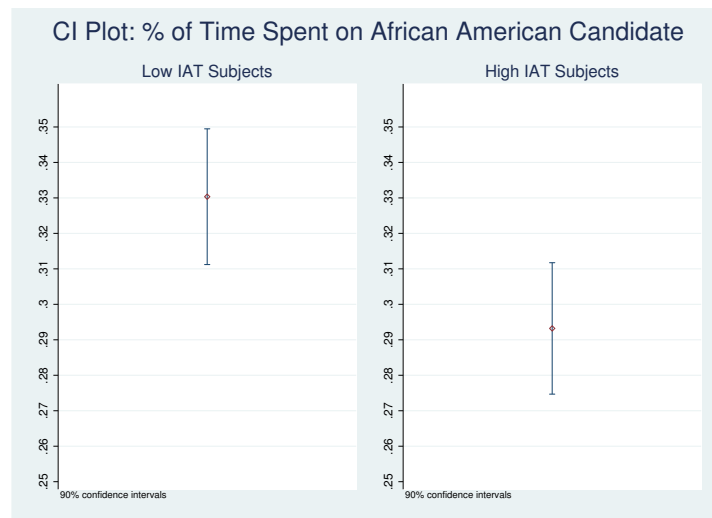
A Mann-Whitney difference of means test shows a statistically significant difference at the .05 level

Figure 4.10: Plotting the effect of IAT across political knowledge

Computed with 90% confidence intervals
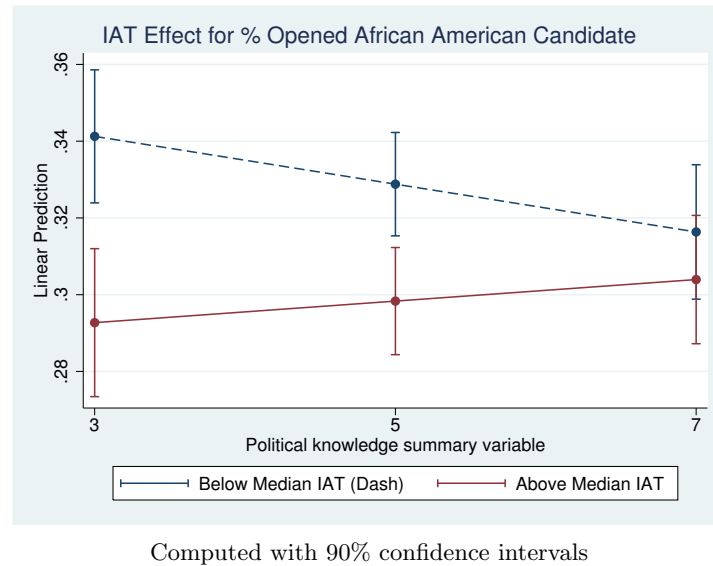
Figure 4.11: Density plot for feeling thermometer difference score across IAT
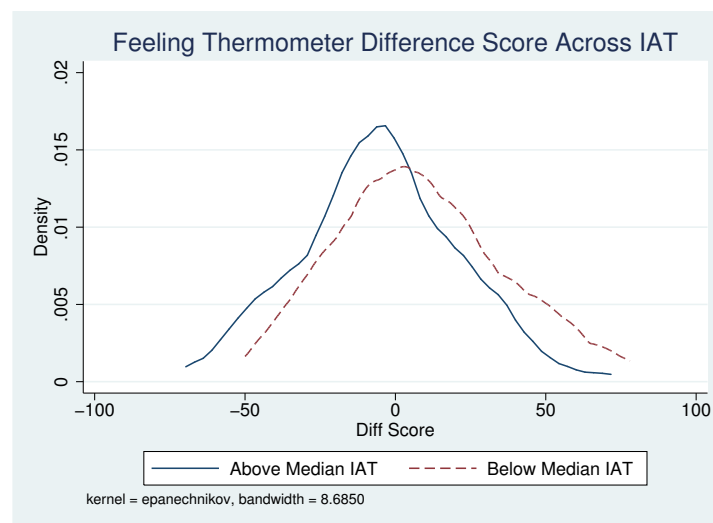
A Kolmogorov-Smirnov test for equality of distribution functions is rejected at the .05 level

Figure 4.12: The Effect of Memory Recall on Candidate Evaluation, Low IAT
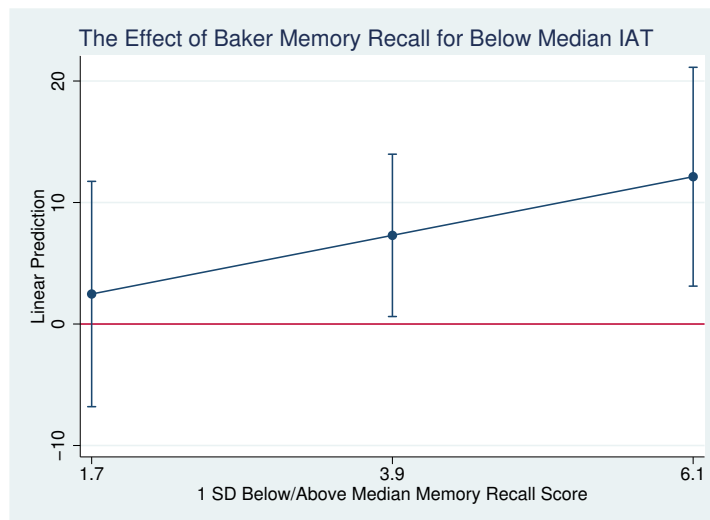
Figure 4.13: The Effect of Memory Recall on Candidate Evaluation, High IAT
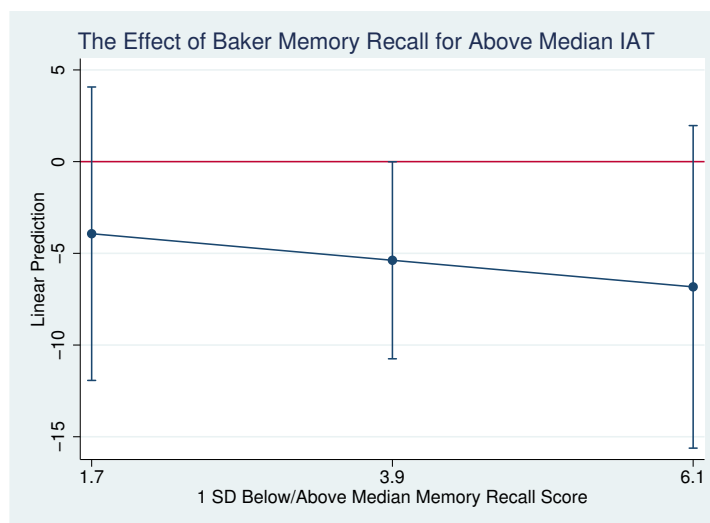
Chapter 5

Implicit Racial Bias and Political Behavior: The Moderating Effect of the Information Environment

5.1 Introduction

According to the dual process literature, automatic or "implicit" attitudes are likely to have a larger impact on behavior in situations where the time and resources to consciously deliberate are limited (Evans, 2003; Payne, 2001; Stanovich and West, 2000; Payne, 2001; Westling, Mann, and Ward, 2006; Shiv and Fedorkhin, 1999). In contrast, deliberative or "explicit" attitudes are likely to be a more significant predictor of behavior when constraints on information and decision-making are relaxed (Fazio and Towles-Schwen, 1999; Strack and Deutsch, 2004). Empirically, there is some evidence to suggest that implicit attitudes' predictive strength is moderated by cognitive capacity. In several studies (Payne et al., 2001; Gibson et al., 2008; Friese et al., 2008), researchers have manipulated subjects' cognitive capacity (via a cognitive load task) and found that implicit attitudes were a significant predictor of choice behavior when cognitive resources were taxed.

However, to date there has been no such analysis examining the impact of cognitive capacity on implicit prejudice as it relates to political decision-making. There are several reasons to think that the interaction between implicit attitudes and cognitive capacity may be particularly relevant for political decision-making. It is well-established that voters rely on a number of simplifying heuristics to navigate the complex and over-abundant amount of information about candidates present in political campaigns (Lau and Redlawsk, 1997; 2006; Bartels, 2000; Lodge and Taber, 2000). This "information-overload" quality that typically characterizes political campaigns may also help to facilitate the activation and use of implicit attitudes to guide voter behavior. Thus, the

impact of implicit attitudes may be particularly strong in context of a campaign where voters are trying to search for information and evaluate candidates under significant informational constraints.

This is coupled by the fact that few studies (if any) have explored the relationship between implicit attitudes and cognitive capacity in the context of racial bias. This may have potentially important implications given that the effect of implicit racial attitudes' may operate differently than other types of implicit attitudes (Fazio and Towles-Schwen, 1997; Dovidio et al., 1997). Specifically, a number of studies have identified conditions under which individuals attempt to override socially sensitive implicit biases. However, as Fazio and colleagues (Fazio and Olson, 1995; Fazio and Towles-Schwen, 1997; Fazio, 2003) point out, such "over-riding" of implicit biases tend to be conditional on how motivated individuals are to suppress such biases.

Hence, the goal of this paper is to empirically examine the potential moderating effect of cognitive capacity on the relationship between implicit racial bias and voter decision-making. Specifically, I use the dynamic information board methodology (Lau and Redlawsk, 1997; 2006) to construct a mock presidential primary campaign where subjects are presented a matrix of "learnable" information about fictional political candidates, which continuously flow down their computer screen for a limited length of time. To test whether the impact of implicit attitudes on political information processing is conditional on cognitive constraints, subjects are randomly assigned to a treatment condition that manipulates cognitive capacity via a number memorization task that subjects must complete throughout the mock campaign and election. At the end of the campaign, subjects are then asked for their vote choice followed by a series of candidate evaluation items. This is followed by the administration of the "black-white Implicit Association Test (IAT) as a measure of implicit racial bias.

This experimental design allows us to assess a number of conjectures related to the potential conditional effects of implicit (and explicit) racial attitudes on political be-havior. Most importantly, I can test whether decreased cognitive capacity facilitates

the activation and use of implicit racial bias, while inhibiting the impact of more deliberative or explicit attitudes. Conversely, I am also able to examine whether increasing cognitive capacity produces the opposite effect, with the predictive strength of implicit attitudes weakening at the expense of explicit ones.

In terms of results, contrary to expectations, I fail to recover a significant IAT effect for any of my behavioral measures in the low cognitive capacity condition. Instead, I find that the explicit prejudice measure is a far stronger predictor of behavior when cognitive capacity is taxed. This is in stark contrast to the high cognitive capacity condition, where implicit racial biases significantly predicts most of my behavioral measures while explicit bias maintains no such effect. Upon first glance, these results appear to contradict the conjectures of dual process theory as to when implicit and explicit attitudes are likely to have their strongest effects.

However, as I detail in great length later on in this paper, these unexpected results are driven primarily by an unintended consequence from the cognitive capacity manipulation. Specifically, I find that rather than overload subjects with information, the treatment condition unintentionally led to subjects exhibiting higher levels of attention and engagement with the experimental tasks, thus crowding out much of the effect of implicit attitudes. Among other indicators, subjects in the treatment condition spent more time reading information about each of the political candidates and recalled more information than those subjects in the control.

Reassessing the results with this in mind offers important insight into the potential conditional effects of implicit racial bias. That increased attention and engagement with the decision task helped to crowd out the influence of implicit attitudes is consistent with the conjecture that explicit attitudes are better predictors of behavior when subjects are sufficiently motivated to engage in deliberative thought (Fazio and Towles-Schwen, 1997). The results from this study have several important implications. Perhaps most importantly, I show that the activation of implicit and explicit biases operate under different sets of conditions. Specifically, individuals are less (more) likely to rely on

their implicit attitudes when they have the (lack) motivation to engage in deliberative behavior. Such a finding has important extensions to voter decision-making. Mainly, implicit racial biases may impact voter decision-making differently for different subsets of the voting population. In particular, these results suggest that implicit racial biases are likely to have a greater influence on those voters that are apathetic or largely disinterested in politics. However, for those motivated voters, implicit racial biases are likely to have little effect.

5.2   Literature Review

Researchers have convincingly demonstrated that implicit attitudes can be important for political decision-making and may impact a variety of important political behaviors (Lodge and Taber, 2008, 2013; Malhotra et al., 2013; Nosek et al., 2009; Payne et al., 2009; Perez, 2010). However, there are number of important outstanding questions that need further exploration. One such question that has received increased attention in recent years is the potential conditional nature of implicit attitudes' effect on behavior. While there is some research on this topic in social psychology and marketing (Fazio and Towles-Schwen, 1997; Friese et al., 2009; Payne et al., 2001), this issue has largely been ignored by political scientists. This has led to a number of critiques of the implicit attitude research program in political science. In their discussion on implicit racial bias, both Huddy and Feldman (2009) and Mitchell and Tetlock (2009) argue that little is known about interactive effects and under what conditions implicit attitudes are likely to maintain their strongest (and weakest) effects. Perez (2013) echoes this sentiment by arguing that one of the new frontiers of implicit attitude research in political science will be to identify both the theoretical and methodological boundaries of implicit attitudes, especially as they pertain to politics.

Dual process theory provides a useful starting point for thinking about potential moderators of implicit attitudes (Bargh, 1992; Bargh, 1999; De Houwer and Moors, 2005; Evans, 2003; Evans, 2010). As mentioned earlier, dual process theory contends

that individuals possess two processing systems, one of which is automatic or implicit (generally referred to as system 1 processing) and the other characterized as more effortful and deliberative (generally referred to as system 2 processing). With regards to the conditions under which individuals are more likely to rely on system 1 rather than system 2 processing, the general consensus in the dual process theory literature is that automatic attitudes have more of an effect in situations where cognitive resources are taxed (Evans, 2003; Payne, 2001; Stanovich and West, 2000). Similar work has been done by Fazio and colleagues (Fazio and Towles-Schwen, 1999; Fazio and Olson, 2003) with their conceptualization of the MODE (Motivation and Opportunity as DEterminants of the attitude-behavior relation) model of behavior. The central tenet of the MODE model is that behavior is driven by controlled processes (i.e. system 2 processing) only when an individual is sufficiently motivated to engage in deliberative thought and has the cognitive capacity to do so.

Empirically, there have been a few studies outside of political science that have started to test these propositions. Payne (2001) examines the strength of affective evaluations in light of time pressure. In his experimental design, subjects participated in a priming task where photographs of White and Black male faces were used as primes, which were then followed by targets that were either handguns or hand tools. As predicted, Payne found that subjects were more likely to incorrectly identify a tool as a gun when the target was primed with a Black face. To test for the effect of time pressure, subjects in the treatment group were randomly assigned a 500ms time deadline to respond to the target stimulus. Payne found that those given the time deadline were more likely to be impacted by the Black photograph prime.

In a second study, Friese et al (2008) examined the predictive validity of implicit attitudes pertaining to consumer behavior when varying subjects' cognitive capacity. Specifically, subjects were first administered both an implicit and explicit measure of their preference toward fruit and chocolate. They were then given a choice task where they were asked to choose five from a selection of ten fruit and ten chocolate items. To

manipulate cognitive capacity, subjects were randomly assigned to one of two experimental conditions. In one condition, subjects memorized a one-digit number during the choice task (i.e. high cognitive capacity). In the other condition, subjects memorized an eight-digit number (i.e. low cognitive capacity) during the choice task. In terms of results, the authors report a significant interaction between cognitive capacity and their implicit measure (but no such interaction with the explicit measure). Hence, the authors conclude that the predictive validity of implicit measures does indeed increase when cognitive capacity is more limited.

5.2.1 Moderators of implicit attitudes in political decision-making

To date, there has been little empirical inquiry into potential moderators of implicit attitudes as it relates to political decision-making and voter behavior. This is somewhat surprising given that informational constraints are often a key feature of political decision-making. As Lau and Redlawsk (2006) note, during an election individuals often encounter an overwhelming amount of information and must deal with cognitive limitations. In one of the few studies to explore potential moderators of implicit attitudes in political science, Kam (2007) shows that the impact of implicit bias against Latinos on preferences for an hispanic candidate is conditional on the absence of a party cue. In Kam's design, subjects are randomly assigned to a treatment condition where subjects receive information pertaining to the party identification of the candidates. Kam finds that implicit (and explicit) bias against Latinos is a significant predictor of vote choice only in the control condition (i.e. no party cue). When party cues are given, the effect of implicit bias appears to disappear. Overall, this suggests that party identification may help individuals override their implicit bias.

In a more recent study, Mo (2011) attempts to model the interaction between implicit bias and partisan identification in the context of the 2008 U.S. presidential election. Mo argues that those who identify as strong partisans possess a stronger motivation to vote for the party candidate and as such should be able override any automatic

predisposition against a candidate due to race. Similar to Kam (2007), Mo finds a statistically significant positive coefficient for the interaction between her implicit racial attitude measure and partisan strength. In her plot of the marginal effect of this interaction, Mo shows that for those with a strong party affiliation (either those that consider themselves a strong Democrat or strong Republican), the impact of implicit bias on vote choice is minimal at best. Conversely, for those who do not have strong party affiliation, implicit bias maintains a strong negative effect on the likelihood of voting for Obama.

While both Kam (2007) and Mo (2011) provide useful starting points for addressing some of the potential moderators of implicit attitudes, it remains unclear the mechanism that is driving the overriding of implicit attitudes in this context, particularly since any number of correlates of partisanship may be responsible for the observed effect (i.e. partisan identification is not randomly assigned). Moreover, the absence or presence of partisan identification may actually reflect the extent of informational constraints present (which would be in support of the hypothesis presented in this study), however this is difficult to disentangle from other potential mechanisms that may arise from cueing party identification. In order to unpack the moderators of implicit attitudes, a cleaner experimental design is needed that allows us to better isolate the potential mechanism(s) at work. The experimental design offered in this study provides such a research strategy. By manipulating cognitive capacity, I am able to cleanly provide a direct test of potential moderators of implicit attitudes that is theoretically grounded in a significant body of literature in cognitive psychology on dual cognition.

5.3 Experimental Design

5.3.1 Participants

This study was conducted in the fall of 2013. All of the experimental sessions were administered in the same laboratory in a large research university. A total of 118 students (52 male and 66 female) were recruited to participate in this study. This study

followed exactly the same procedures as study one (in Chapter 4). Approximately one week prior to the lab sessions, subjects were administered a pre-study questionnaire, which included demographic items and two measures of explicit racial prejudice. The demographic items included gender, race, and major. For the explicit racial prejudice measure, I use Kam and Kinder's (2008) ethnocentrism scale as well as Kinder and Sear's (1982) symbolic racism scale. Finally, I also include a measure of subjects' political ideology and party identification.[1]

Upon arrival to the lab, subjects were given a short political knowledge quiz before starting the mock presidential primary campaign. The political knowledge quiz consists of eight items and covers a variety of topics including questions on the branches of government, identifying well-known political figures, and foreign policy.[2] Once subjects completed the mock primary campaign and voted in the election, they moved to the final part of the study, which is the administration of the black-white IAT. Again, the mock primary campaign and black-white IAT were exactly the same versions that were administered in the previous study from chapter 4. However, in this second study, I modify the experimental design by adding the cognitive load manipulation. Before turning to results, I will first provide some more detail on how this manipulation worked.

5.3.2 Cognitive load manipulation

Cognitive capacity or cognitive load refers to the use of working memory to process and learn information (i.e. to form schemas that link information processed in the working memory to the long-term memory) (Paas et al., 2004). Since working memory capacity is extremely limited (see Chapter 2), the amount of information that individuals can process simultaneously can either under-load or overload the amount of working memory available. If working memory is sufficiently overloaded, then information is not processed and thus will not be learned. The most frequently used method to manipulate cognitive load is to have subjects memorize strings of numbers while they are

---

[1]The text for all of these items can be found in the Chapter 3 appendix.

[2]See the Chapter 3 appendix for a list of the questions.

simultaneously performing another task (Swann et al., 1990; Shiv and Fedorikhin, 1999; Friese et al., 2008).

Following previous studies, subjects in study two are randomly assigned to a high-load treatment group.[3] The high load treatment condition follows the same procedures as study one and does not differ in content (i.e. the candidates and the information available to subjects are exactly the same as study one). Rather, the only difference is the inclusion of a treatment, which seeks to manipulate subjects' cognitive load by asking them to memorize (and recall) sets of five-digit number strings while they participate in the mock presidential primary election.[4] Specifically, subjects are randomly assigned into one of three sub-conditions. In one sub-condition subjects are given two sets of number strings to memorize (and recall) during the campaign, but are not asked to memorize any number strings during the candidate evaluation/vote choice phase.[5] In a second sub-condition, subjects are only given one number string to memorize (and recall) at the onset of the candidate evaluation/vote choice phase.[6] Finally, in a third sub-condition, subjects are given the memorization task both during the campaign and during the evaluation period.[7] In order to assess whether the cognitive resources manipulation had its intended effect, I use (as a manipulation check) the memory recall task given to subjects at the end of the election. Under the assumption that high cognitive load impairs individuals' ability to learn new information, I expect that memory recall should be lower in each of the high load conditions (as compared to study 1, which had no cognitive load manipulation).

---

[3]As will be discussed in section 4.2.1, study 1 (chapter 4) and study 2 (i.e. this study) were administered during the same administration window and subjects were randomly assigned to one of the two studies. The results from study 1 and study 2 will be compared in subsequent sections of this paper.

[4]For the exact wording of this manipulation, please see the Chapter 3 appendix.

[5]Specifically, one number string is given right before the start of the primary and the recall question is approximately seven minutes later. A second number string is given in the ninth minute of the primary and the recall question is again asked seven minutes later.

[6]The number string is administered immediately after the information search and the recall question is asked after the vote choice, feeling thermometers, and memory recall tasks are finished.

[7]In this condition, a total of three number strings are administered. Two during the information search (approximately seven minutes apart) and one at onset of the candidate evaluation stage

The decision to include three sub-conditions for the cognitive load manipulation stems from the fact that it is not entirely clear how the cognitive load manipulation will operate. Previous studies that have used the memorization task to manipulate cognitive load typically include a choice task that takes no longer than one minute. Since the choice task is quite short, previous studies only use one set of number strings to manipulate cognitive load. In this study, there are two sets of choice tasks: one during the primary campaign and another when subjects must evaluate the candidates. During the primary, the choice tasks consists of what information to look at for each candidate. Unlike previous studies, this choice task is continuous and quite long (approximately sixteen minutes). Hence, it is unclear as to whether the memorization of one number string during the entirety of the mock campaign (and election) is the most effective way to manipulate cognitive load. To account for this uncertainty, I use three different manipulations that vary the number of memorization tasks (and time between recall). This will allow us to explore whether differences exist on account of the number of memorization tasks given to subjects. In subsequent analysis not shown (but available upon request), I find no significant differences between my sub-conditions, hence the data from each sub-condition is pooled in order to produce one treatment condition.

5.4   Results

5.4.1   Descriptive analyses

Similar to study one, I can start my analysis by first focusing on the information search variables. Figure 5.1 displays the density plot for the percentage of information opened for the African American and white moderate candidates.[8]

- Insert Figure 5.1 Here -

Like study 1, I see very little difference in the distribution of opened items for the two candidates. Does this change when I condition on IAT score? Figure 5.2 shows

---

[8]Again, recall that due to randomization, these candidates are in expectation equivalent.

a density plot for the percentage of opened items for the African American candidate now conditioning on the IAT median split variable.

- Insert Figure 5.2 Here -

Unlike study 1, where I reported large (and statistically significant) differences in both the distribution of opened items and point estimates between the high and low IAT groups, the results from study 2 show no such differences. Both the Kolmogorov-Smirnov test for equality of distributions and the Mann-Whitney difference of means test fails to be rejected at the .10 level.

Turning to the percentage of time spent reading information about the African American candidate, I again observe the same pattern. Figure 5.3 shows the density plot of the duration variable for the African American and white moderate candidates while Figure 5.4 shows the density plot for the African American candidate now conditioning on IAT score. I again fail to recover any difference in reading time for the low and high IAT groups.

- Insert Figures 5.3 and 5.4 Here -

In addition to the information search variables, I can also explore whether there are differences between our IAT groups with respect to the candidate evaluation measures. Recall that in study one, high IAT subjects rated the white moderate candidate higher than its African American counterpart compared to the low IAT group and this difference was statistically significant. Figure 5.5 presents the density plots of the feeling thermometer difference score for the high and low IAT groups. Again I see no difference in the distribution of scores for my two groups. Moreover, a difference of means test also fails to show any statistically significant differences.

- Insert Figure 5.5 Here -

In sum, upon first glance IAT appears to have no effect on information search in the treatment condition. Specifically, within the high cognitive load treatment group, I find very little difference between the high and low IAT groups across each of my behavioral outcomes. However, missing from this analysis is whether the explicit prejudice measure experienced the same moderation. In other words, is it the case that the cognitive load manipulation also reduced the predictive strength of explicit racial bias? To address this question, I can turn to the multivariate analyses for both the information search and candidate evaluation measures.

5.4.2   Regression analysis

The focus of this section is to examine the impact of my cognitive load treatment on political information search and voter decision-making.[9] To assess my treatment effect, I want to compare subjects' behavior under high and low cognitive capacity. To do this, I can pool the data from study 1 (i.e. no cognitive capacity manipulation) and the present study (i.e. study 2) where I manipulate cognitive capacity (see Chapter 3 for more details). Pooling the data in this way also makes sense from a methodological standpoint since study 1 and study 2 are exactly equivalent but for the cognitive capacity manipulation. Moreover, both studies were conducted during the same administration window with subjects being randomly assigned to one of the two studies. Hence, in expectation there should be no differences between the two samples.[10] For ease of reporting the results, I will refer to study 1 as the "control" condition and study 2 as the "treatment condition."

I start my multivariate analysis by examining the effect of the treatment on both implicit and explicit prejudice. For the first set of models, I assess the treatment effect on the total percentage of opened items for the African American candidate. Table 5.1 displays these results.

---

[9]For all analyses, I use robust standard errors. Alternative models without robust SEs yield no substantive difference. Results from these models are available upon request.

[10]Please see the next section for descriptive analyses confirming balance on observables between the two studies.

- Insert Table 5.1 Here -

For the first model, I simply show the results when I pool all of the data and omit my controls. Here I see that both the implicit and explicit measures go in the expected direction, however only the explicit prejudice measure is significant. In model 2, I interact a treatment dummy variable with both the implicit and explicit measures and again I omit the control variables. The results from this model suggests potential conditional effects for both racial measures. However, the question remains whether such interaction effects are robust to the inclusion of the controls. Model 3 displays the same set of interactions but this time including all of the control variables. Again, I find significant interaction effects for both of the implicit and explicit measures, which suggests that the slope of the IAT effect is different for the control and treatment groups. To better understand the nature of these relationships, I plot the interaction effects for both the implicit and explicit measure.

- Insert Figures 5.6 and 5.7 Here -

In figure 5.6, I graph the slopes for IAT's effect (on the percent opened dependent variable) for the control and treatment conditions. I see that the slope for the control condition (i.e. no cognitive capacity manipulation) is negative and statistically significant (at all values above the median IAT score) when compared to the slope of the treatment condition. This provides further evidence to suggest that the treatment condition significantly weakened the effect of implicit prejudice. Turning to figure 5.7, I display the same plot but now for the explicit prejudice measure. Here I find the exact opposite effect with the slope for the treatment condition being negative and statistically significant (at most values above the median score) when compared to the control condition. Hence, it appears that the treatment condition actually strengthened the predictive power of explicit prejudice.

Turning to the reading time dependent variable, I observe the same pattern of effects.

Table 5.2 provides the results from the same three regression models used for the percentage opened variable.

- Insert Table 5.2 Here -

Again I recover significant interaction effects between the treatment dummy and both prejudice measures. Figures 5.8 and 5.9 plot the slopes for these interaction effects and again I report the same set of findings as I previously observed in figures 5.6 and 5.7.

- Insert Figures 5.8 and 5.9 Here -

Hence, for the information processing variables, it appears that the cognitive load treatment is associated with an increase (decrease) in the activation and use of explicit (implicit) prejudice to help guide information search. To examine whether the reported treatment effect is restricted to information processing, I can also explore potential interaction effects for the candidate evaluation measures. Tables 5.3 and 5.4 display the results when I regress the feeling thermometer difference score and the vote choice measure on the interaction between the implicit and explicit measures and the treatment dummy.

- Insert Tables 5.3 and 5.4 Here -

In both sets of results, the IAT effect differs by experimental condition. Specifically, IAT's effect is statistically significant in the control condition but loses significance in the treatment group. Turning to explicit prejudice, I observe findings consistent with the other models in that explicit prejudice appears to have a larger negative slope in the treatment condition (relative to the control), however these effects are weaker and fail to reach statistical significance.

In total, the findings from the comparison of my first and second study appear to be almost entirely at odds with my original hypotheses. That automatic attitude activation

appears to be facilitated by increased cognitive capacity while explicit attitudes are more likely activated under low cognitive capacity is directly at odds with the tenets of dual process theory. These counterintuitive results raise several important questions pertaining to the effectiveness of my experimental design. In the next section, I will spend some detail attempting to shed light as to what can account for these surprising results.

5.5 What did the manipulation actually do?

Before closely examining some of the results from my manipulation checks, it is important to first analyze whether balance was achieved across observables for the two conditions. Recall that study 1 (i.e. the "control" condition) and study 2 (i.e. the "treatment" condition) were conducted during the same administration window and when subjects showed up to the lab they were randomly assigned to one of the two studies. Hence, in expectation, there should be no significant differences between the samples of the two studies. To ensure that this was indeed the case, a balance check was administered for all of observables. The results confirm that there are no significant differences in observables across the two conditions[11]

By ruling out potential sample imbalances, I can turn my attention to exploring how the manipulation actually operated. As mentioned earlier in this paper, researchers often attempt to manipulate cognitive load via a memorization task that is administered while subjects engage in the choice task. Since individuals' working memory capacity is severely limited (Miller, 1956), the memorization task is meant to overload subjects with information thus making it difficult to process new information and helping to facilitate dominant responses (Banaji and Greenwald, 1995). However, unlike previous studies using this procedure, I have built into the study a series of manipulation checks to examine whether the cognitive load treatment actually had its intended effect on information processing capacity.

---

[11]Please see the appendix for these results.

One set of manipulation checks is to compare the information search variables across my two conditions. To the extent that the cognitive load manipulation impaired subsequent information processing, I should see significant differences in both the total number of opened items and the total reading time across the two conditions. Specifically, one would expect that in the cognitive load condition, subjects will have opened less items (in total) and have longer reading times (in total) than compared to the control condition. Figures 5.10 and 5.11 compare the density distributions for total opened and total reading time across conditions.

- Insert Figures 5.10 and 5.11 Here -

For the total opened variable, I surprisingly see no difference in the distribution functions across the conditions. Moving to the comparison of total duration, I again observe no statistically significant differences between the treatment and control conditions. Hence, contrary to expectations, the manipulation did not appear to induce the "informational overload" quality that was intended. However, from these initial results it remains difficult to decipher the effect of the manipulation on processing capacity, if there was an effect at at all. To do this, I need to dig deeper into the data.

Another information processing variable that should provide leverage in understanding the treatment effect is memory recall. Again, if the manipulation operated as intended, then one would expect that in the cognitive load condition, subjects should have a more difficult time recalling information about each of the candidates. In fact, the negative impact of cognitive load on memory recall is the central (often untested) assumption made by researchers attempting to manipulate processing capacity. Fortunately, by having a memory recall task at the end of the mock primary, I am in a position to directly assess this relationship. Specifically, I compare the total memory recall (i.e. summing the number of information items recalled for each candidate) across the conditions. In the control condition, I find that the mean number of information items recalled is 12.2. Compare this to the treatment condition, where the mean number

is 13.4. This difference is statistically significant at the .10 level. Once again, counter to expectations, I find that those in the cognitive load condition were actually able to recall more information than those in the control. Such findings suggest that rather than manipulate cognitive capacity, the treatment may have had the unintended effect of increasing subject's engagement/motivation with the task.

Continuing with the examination of the relationship between the cognitive load condition and memory recall, an alternative manipulation check is to examine the association between memory recall and information search behavior. In other words, to the extent that subjects are motivated or engaged with the task, this should be reflected in the strength of association between the percentage of opened items for any given candidate and the amount of information subjects are able to recall. The assumption here being that, on average, the more information items a subject opens about a given candidate, the more pieces of unique information he/she should be able to recall about said candidate. Figures 5.12 and 5.13 present scatterplots for the percentage of opened items for the African American candidate and the number of information items recalled for the African American candidate in the memory task across the treatment and control conditions.

- Insert Figures 5.12 and 5.13 Here -

I find that the correlation between the percentage of opened items and memory recall is much stronger in the treatment condition compared to the control. This is further evidenced by comparing the Spearman correlation coefficients, which is .24 in the control compared to .44 in the treatment condition. Moreover, these findings are not restricted to the African American candidate. In fact, I observe the same pattern in the data for both the white moderate candidate as well as the partisan candidate. Again, these results are consistent with the story that the experimental manipulation induced heightened engagement/motivation levels among subjects.

A potential objection to these findings is that I have yet to rule out whether for some idiosyncratic (and unobservable) reason, subjects in the treatment condition just happened to be a more attentive bunch than the subjects in the control. If that is the case, then of course this confounds the hypothesized treatment effect. To ensure that this is not the case I can examine the time spent reading the instructions that subjects were given at the onset of the experiment. Specifically, I can compare the reading time across conditions for the set of instructions describing the dynamic process tracing environment. Comparing the reading time for these instructions is a plausible proxy for subjects' initial motivation/engagement since all subjects are unfamiliar with the dynamic process tracing environment; hence, the expectation should be that all subjects should carefully read these set of instructions.

In terms of details for the instructions, prior to start of the primary campaign, subjects were given a screen of instructions that described how to navigate the dynamic process tracing environment. Following these instructions, subjects were given a three minute demonstration version of the task, which did not include the candidates used for the actual primary. Since these instructions were exactly the same across conditions and were given prior to the cognitive load manipulation, they provide an excellent opportunity to assess whether there were any pre-manipulation differences in engagement/motivation between conditions. Table 5.5 shows the median reading time for the practice instructions. There is virtually no difference between the conditions, thus adding more evidence to suggest that the observed differences between conditions was not driven by unobservable difference in engagement/motivation across samples.

- Insert Table 5.5 Here -

To sum up, after assessing the effect of the manipulation, I conclude that rather than decreasing cognitive capacity, the treatment condition instead primed subjects to be more motivated and engaged with the task. Subjects in the treatment condition spent more time reading information about each of the candidates, they were better

able to recall information about each of the candidates, and their reported correlations between the percentage of opened items and memory recall were far higher. Moreover, I find no differences in observables that may be driving such effects and my examination of the pre-manipulation reading times for the experimental instructions yields no significant differences across conditions thus further demonstrating that there were likely no unobserved differences in motivation/engagement between the two groups.

The question remains as to why the cognitive capacity manipulation yielded such drastically different effects as compared to previous studies that used the same number memorization task. One hypothesis relates to the nature of the decision task. As mentioned earlier, previous studies examining the effect of cognitive capacity on decision-making use a choice task that is quite short. As such, subjects typically do not memorize more than one set of five digit numbers. In contrast, this study employed a decision task that was quite long (approximately sixteen minutes) and typically required subjects to memorize (and recall) more than one set of numbers. Such differences may have had the unintended effect of priming subjects to pay closer attention throughout the duration of the decision task, which inevitably led to longer reading times and better memory recall. A second (but related hypothesis) relates to the priming of a potential accuracy bias. As part of the manipulation, by emphasizing that subjects accurately recall the five-digit number strings, the treatment condition may have also primed subjects to treat the primary campaign itself as an accuracy task. As a result, subjects may have engaged in more active information search throughout the campaign. However, subsequent work will be need to be done in order to empirically assess these hypotheses.

5.6   Implicit bias and the motivation to control prejudice

Returning to the results, my main findings were that implicit racial bias (and not explicit racial bias) was a significant predictor of both information search and candidate evaluation in the control condition; however implicit attitudes had no such effect in

the treatment condition with instead explicit prejudice being the primary predictor of behavior. While these results are seemingly counterintuitive, seen through the lens of subjects' motivation levels these findings appear quite plausible. Fazio's (1990; Fazio and Towles-Schwen, 1999) MODE model (Motivation and Opportunity as DEterminants of the attitude-behavior relation) contends that while implicit attitudes can at times directly influence behavior, it's effects are often mediated by individuals' motivation to engage in deliberative information processing. For instance, several studies show that subjects relied less on implicit attitudes when they were given a treatment that explicitly motivated them to process information carefully (Sanbonmatsu and Fazio, 1990; Schuette and Fazio, 1995).

In this light, this study's findings no longer appear counterintuitive. Instead, the results provide some evidence that individual can override their implicit biases when they are motivated to do so. In this study, by informing subjects that they will be memorizing strings of numbers while participating in the mock election campaign, I appeared to increase motivation during information processing.[12] As a result, the predictive strength of implicit bias dramatically weakened in the treatment condition. Instead, I saw that deliberative attitudes (mainly explicit prejudice) were a far better predictor of behavior.

A natural extension to this research is to explore whether the treatment had differential effects on subjects. Specifically, depending on their levels of implicit and explicit racial bias, some subjects may be more (or less) less likely to engage in the motivated overriding of implicit biases. For instance, Mo (2011) argues that motivated overriding should only occur among those that maintain low levels of explicit racial bias since those that maintain high levels of explicit racial bias (in addition to implicit racial bias) are presumably not motivated to moderate their implicit bias. Since I have data on both

---

[12]As mentioned throughout this paper, it remains unclear as to why the treatment condition had this unintended effect. One hypothesis is that by asking subjects to memorize number strings prior to primary campaign, the manipulation may unintentionally have primed subjects to pay more attention to the stimuli as compared to the control. However, given data limitations, there is no way to empirically test this hypothesis.

implicit and explicit bias in addition to a randomly assigned manipulation that appears to prime motivation, I can test this hypothesis.

To examine this, I can look at the interaction between the treatment dummy and the implicit bias measure for those above and below the median the explicit bias score. In terms of expectations, if Mo (2011) is correct, then I should see the treatment having little effect on those with higher levels of explicit bias and much a larger effect on those lower levels of such bias. Again, the reason we may expect this is because as argued by Mo (2011), those with high levels of explicit bias have little motivation to override their implicit bias since they have no problem explicitly articulating their discriminatory preferences. This is also consistent with the work by Fazio and colleagues (1990, 1995, 1997), which argues that only those individuals sensitive to social desirability will attempt to override their implicit biases.

Table 5.6 provides the results from the interaction model where I regress percentage of opened items for the African American candidate on the standard set of covariates including the interaction between the treatment dummy and IAT scores, however the key here is that I am sub-setting the data to those who scored above the median explicit bias score and those that score at or below the median explicit score.

- Insert Table 5.6 Here -

More importantly, Figures 5.14 and 5.15 plot the marginal effect of the treatment across IAT score subset by high and low levels of explicit bias. What stands out here in the comparison of the two figures is that for high explicit bias individuals, the marginal effect of the treatment is indistinguishable from zero while for the low explicit bias individuals, I see that the marginal effect of the treatment is positive and statistically significant.

- Insert Figures 5.14 and 5.15 Here -

These differences suggest that the effect the treatment on implicit prejudice primarily operates for those with low levels of explicit bias. Such findings are consistent with the story that while motivation can help individuals override their implicit biases, we should only see such effects when subjects are at least outwardly anti-prejudice. For those individuals that are open and explicit about their racial bias, increasing the motivation to override their implicit biases will have little effect.[13]

5.7   Conclusion

Recent research in the implicit attitudes literature has focused on potential moderators of implicit bias. While there is some research on this topic in social psychology and marketing (Fazio and Towles-Schwen, 1997; Friese et al., 2008; Payne et al., 2001), this issue has largely been ignored by political scientists. This has led to a number of critiques of the implicit attitude research program in political science. Moreover, it has also led to some recent studies questioning whether implicit prejudice (or implicit attitudes more generally) has much of an impact on political judgment and decision-making above that of explicitly measured racial attitudes (Ditonto, Lau, Sears, 2013; Ryan and Kinder, 2013; Segura and Valenzuela, 2010).

Hence, the focus of this study was to shed further light on this subject by attempting to manipulate cognitive capacity in ways that would either facilitate or inhibit the activation of implicit attitudes. Specifically, I randomly assigned subjects to a memorization task meant to overload individuals's cognitive capacity in an attempt to facilitate the use of implicit attitudes. However, as I carefully demonstrated, the effect of this study's manipulation operated in an entirely unexpected way with those assigned to the memorization task exhibiting higher levels of motivation and engagement throughout the experiment. Fortunately, this unintended effect still yielded a set of interesting results that provide important insight as to when implicit attitudes are likely to have stronger and weaker effects.

---

[13]Please note that similar models were run for the other behavioral dependent variables and I find very similar effects.

Specifically, my findings show that increased motivation/engagement inhibits the activation of implicit attitudes. While I found implicit racial bias to have a significant effect on both political information search and candidate evaluation in the control condition, I find no such effect for those in the treatment group. Moreover, for explicit prejudice, I observe the exact opposite pattern of results with no statistically significant effects in the control condition but statistically significant negative effects for those in the treatment group.

In sum, this study's findings provide important insights into understanding when implicit attitudes (and specifically implicit racial bias) are likely to be stronger (or weaker) predictors of political behavior. In particular, the results suggest that increased motivation/engagement with the decision task helps override implicit biases. However, I also show that such effects only occur for those with low levels of explicit bias. In other words, if individuals' implicit and explicit preferences are aligned then such overriding is unlikely to occur.

These findings have a number of important implications for the literature on implicit attitudes in political science. Perhaps most importantly, it suggests that the influence of implicit attitudes on decision-making is very much conditional on the nature of the information environment. Implicit biases are likely to have their strongest influence on behavior when individuals are relaxed and rely on their normal routines. Conversely, the effect of implicit attitudes is likely to be attenuated when individuals are in heightened state of awareness. This may help explain why previous findings show implicit attitudes having stronger effects of non-verbal behaviors in relation to explicit attitudes (Dovidio et al., 1997). Moreover, this may also help to explain why there is so much variance in the reported effect size of implicit attitudes on various forms of behavior.

5.8   Tables and Figures

Table 5.1: Information opened for African American candidate, across conditions

	(1)	(2)	(3)
IAT	-0.025	-0.075**	-0.078**
	(0.023)	(0.032)	(0.033)
Explicit Bias	-0.003**	0.0002	0.0001
	(0.001)	(0.001)	(0.002)
Treatment		-0.003	-0.007
		(0.014)	(0.014)
Treatment * IAT		0.098**	0.113**
		(0.047)	(0.068)
Treatment * Explicit		-0.006***	-0.007***
		(0.002)	(0.002)
Controls	No	No	Yes
Constant	0.327***	0.329***	0.343***
	0.007	(0.009)	(0.076)
N	233	233	233
R^2	0.026	0.060	0.117

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]Controls include gender, race, political knowledge, partisanship, and the total number of opened items

Table 5.2: Reading time for African American candidate across conditions

	(1)	(2)	(3)
IAT	-0.030	-0.094**	-0.103**
	(0.032)	(0.047)	(0.049)
Explicit Bias	-0.004**	0.001	0.001
	(0.002)	(0.002)	(0.002)
Treatment		-0.005	-0.010
		(0.018)	(0.019)
Treatment * IAT		0.125**	0.149**
		(0.065)	(0.069)
Treatment * Explicit		-0.009***	-0.010***
		(0.003)	(0.003)
Controls	No	No	Yes
Constant	0.327***	0.330***	0.264
	0.009	(0.013)	(0.220)
N	233	233	233
R^2	0.024	0.061	0.110

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]Controls include gender, race, political knowledge, partisanship, and the total duration and reading speed

Table 5.3: Regression for candidate evaluation, across conditions

	(1)	(2)	(3)
IAT	-11.570	-30.937*	-30.628**
	(8.943)	(16.356)	(15.621)
Explicit Bias	-1.831***	-1.243	-1.005
	(0.638)	(0.995)	(1.020)
Treatment		-7.359	-10.074*
		(5.801)	(5.373)
Treatment * IAT		34.424*	36.136**
		(19.433)	(17.649)
Treatment * Explicit		-1.132	-0.908
		(1.234)	(1.211)
Controls	No	No	Yes
Constant	4.756*	9.131*	8.939†
	2.551	(5.133)	(5.567)
N	233	233	231
R^2	0.050	0.066	0.272

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]Controls include gender, race, political knowledge, partisanship, and memory count

Table 5.4: Vote for African American candidate

	(1)	(2)	(3)
IAT	-1.403†	-2.839*	-3.343**
	(0.881)	(1.506)	(1.742)
Explicit Bias	-0.152**	-0.096	-0.207
	(0.070)	(0.124)	(0.174)
Treatment		0.018	-0.282
		(0.484)	(0.576)
Treatment * IAT		2.491	2.596
		(1.871)	(2.150)
Treatment * Explicit		-0.092	0.071
		(0.152)	(0.202)
Controls	No	No	Yes
Constant	-0.705***	-0.694*	-0.324
	0.230	(0.383)	(1.088)
N	233	233	231
R^2	0.039	0.054	0.274

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]Controls: gender, race, political knowledge, and partisanship
[b]Outcome variable is 1 if voted for African American and 0 otherwise.

Table 5.5: Median Reading Time for Practice Instructions

Condition	Median
Control	65.58
Treatment	65.66
Total	65.58
N	115

Table 5.6: Treatment effect on information opened for low/high explicit bias

	High Explicit Bias	Low Explicit Bias
Female	-0.002	-0.009
	(0.015)	(0.012)
White	-0.038**	-0.022**
	(0.016)	(0.011)
Political Knowledge	0.006✝	0.0001
	(0.004)	(0.002)
Liberal	0.055**	-0.009
	(0.026)	(0.017)
Conservative	0.066**	0.011
	(0.028)	(0.022)
Total Opened	0.005	0.005
	(0.030)	(0.018)
IAT	-0.154***	-0.051
	(0.058)	(0.042)
Treatment	-0.046✝	0.004
	(0.031)	(0.016)
Treatment * IAT	0.118	0.073
	(0.101)	(0.059)
Constant	0.287**	0.328
	(0.135)	(0.091)
N	91	142
R^2	0.167	0.08

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]Controls include gender, race, political knowledge, partisanship, and the total number of opened items

Figure 5.1: Density plot, % of opened items for the two moderate candidates
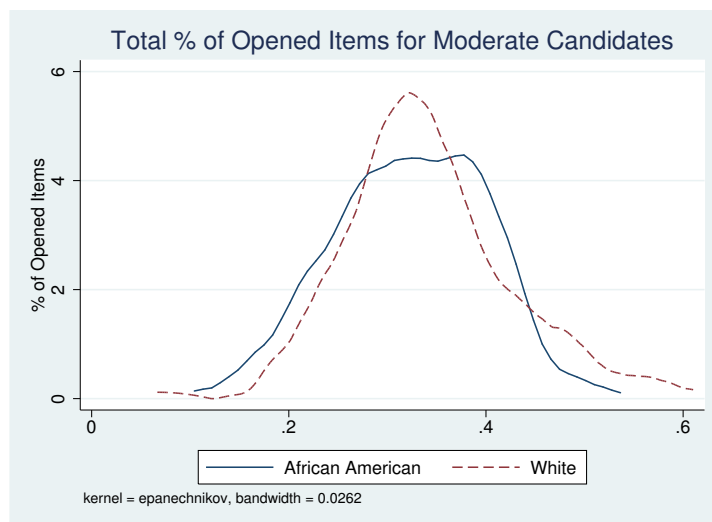
Figure 5.2: Density plot, % of opened items for African American candidate by IAT
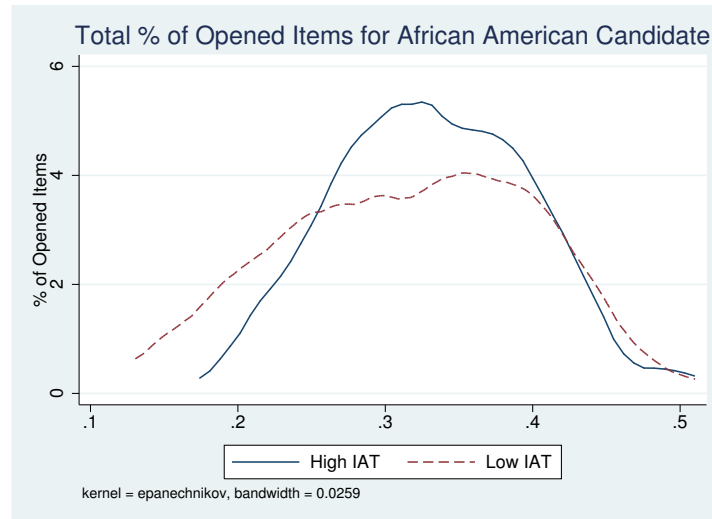
A Kolmogorov-Smirnov test for equality of distribution functions is rejected at the .01 level

Figure 5.3: Density plot, % of reading time for moderate candidates
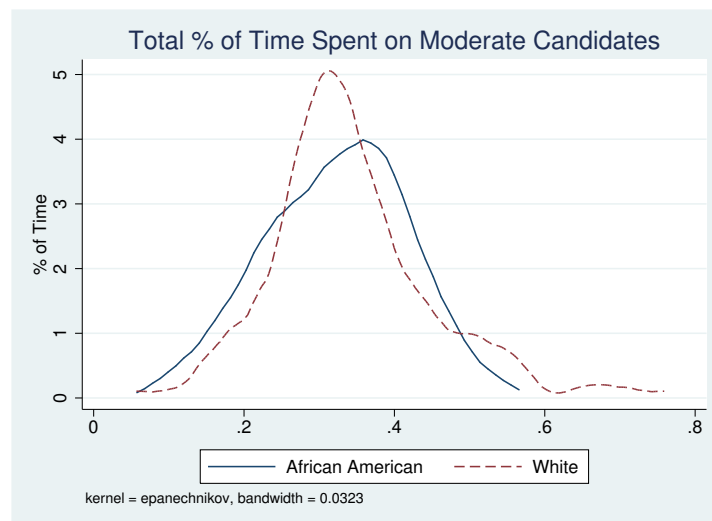
Figure 5.4: Density plot, % of reading time for African American candidate by IAT
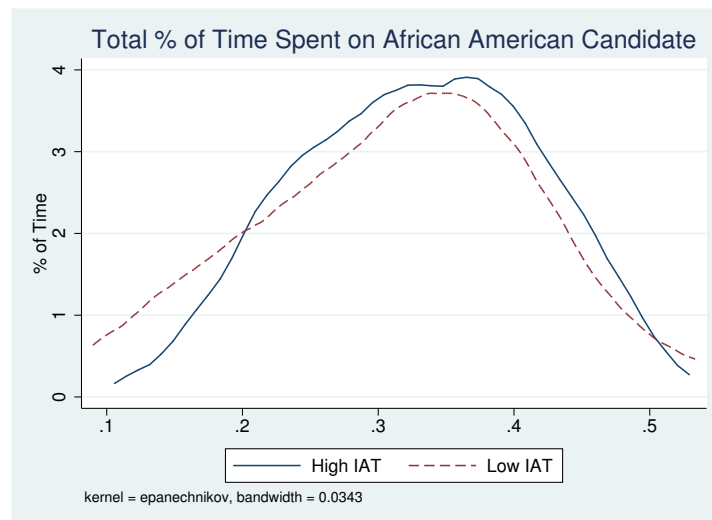
A Kolmogorov-Smirnov test for equality of distribution functions is rejected at the .05 level

Figure 5.5: Density plot, feeling thermometer difference score by IAT
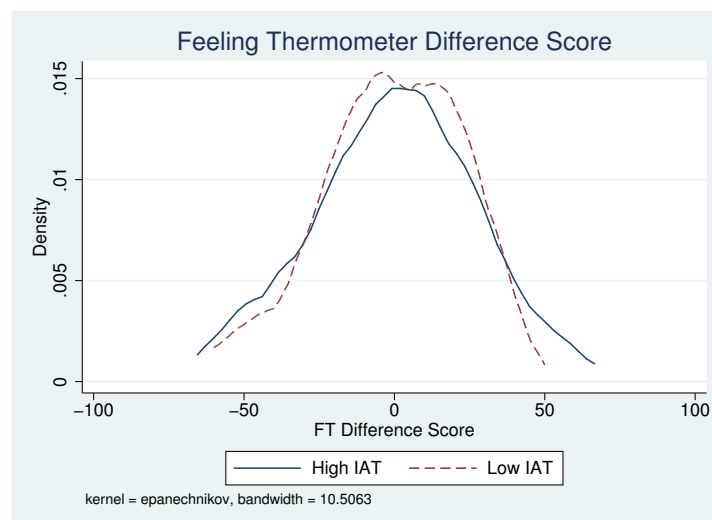
A Kolmogorov-Smirnov test for equality of distribution functions is rejected at the .05 level

Figure 5.6: Interaction plot, implicit prejudice across conditions

% Opened Information for African American Candidate
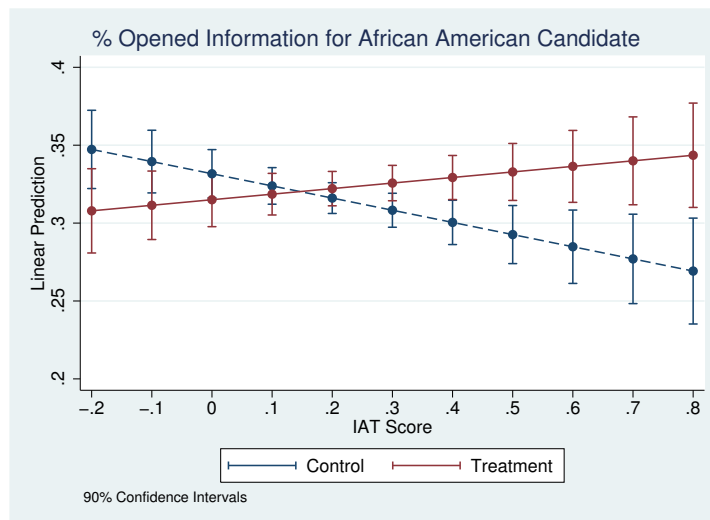
Figure 5.7: Interaction plot, explicit prejudice across conditions

% Opened Information for African American Candidate
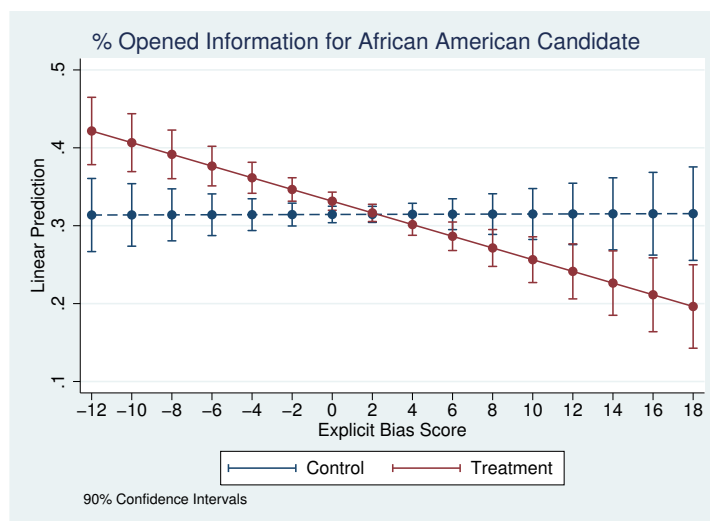
Figure 5.8: Interaction plot, implicit prejudice across conditions
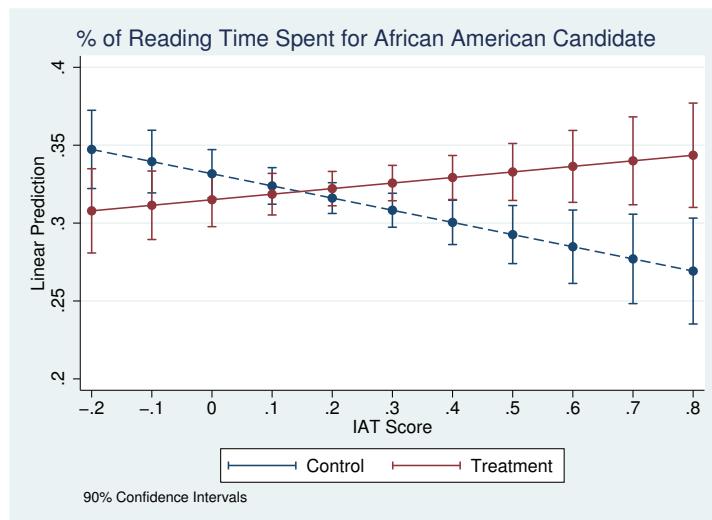
Figure 5.9: Interaction plot, explicit prejudice across conditions
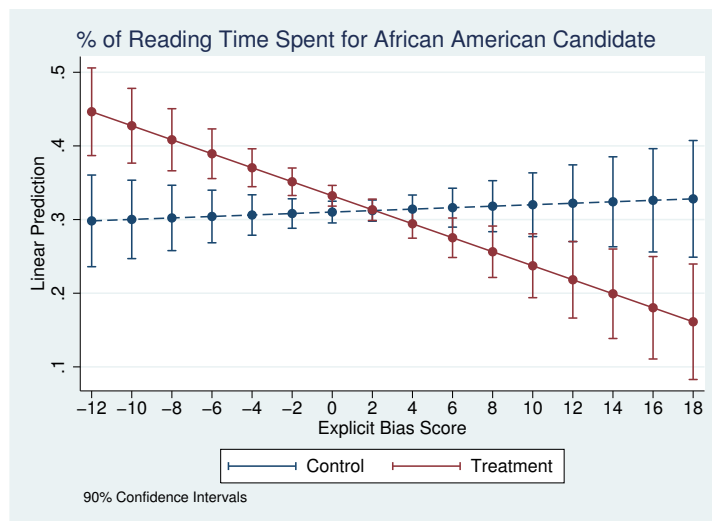
Figure 5.10: Density plot, number of opened items for African American candidate across condition
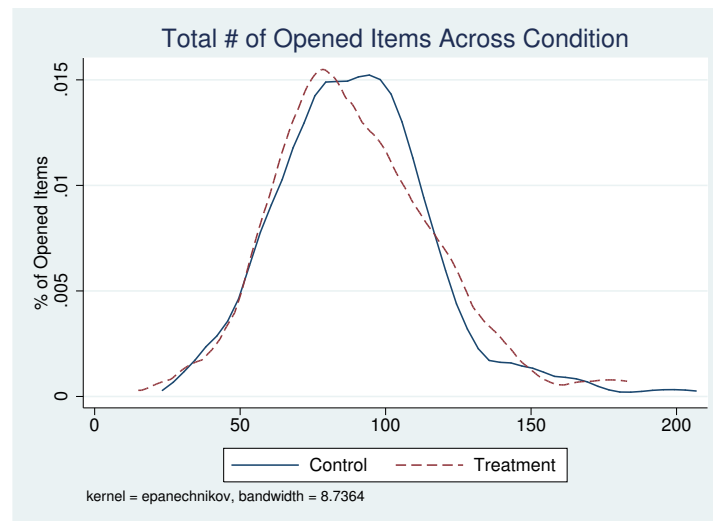
Figure 5.11: Density plot, total reading time for African American candidate across condition
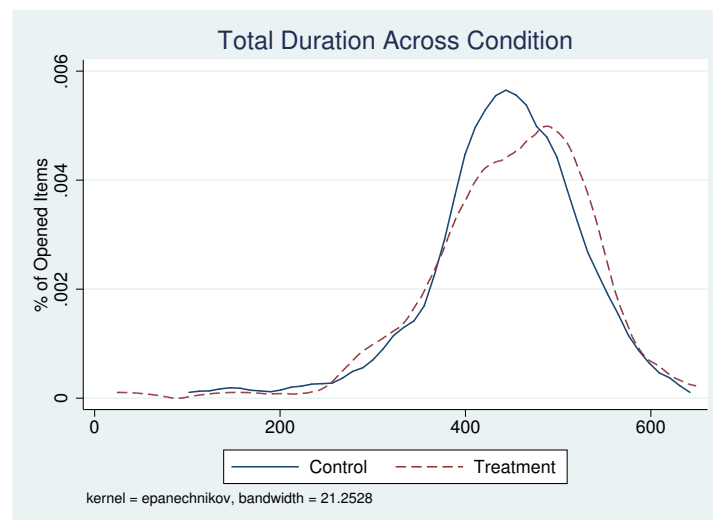
A Kolmogorov-Smirnov test for equality of distribution functions is rejected at the .12 level

Figure 5.12: Association Between Total Opened and Memory Recall: African American candidate
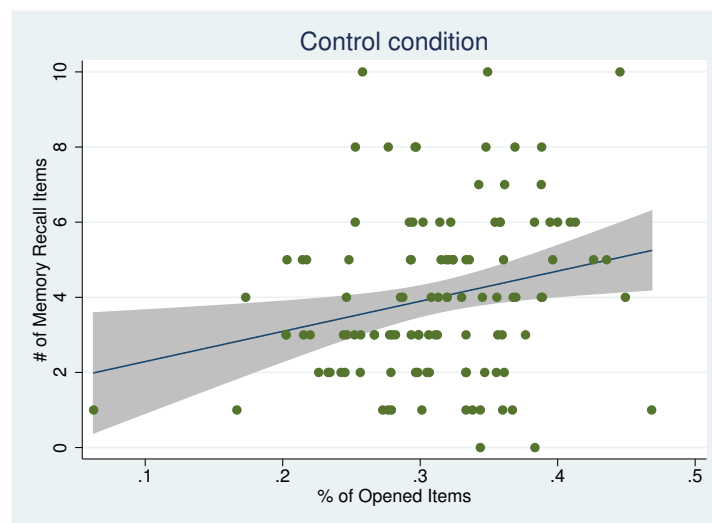
Figure 5.13: Association Between Total Opened and Memory Recall: African American candidate
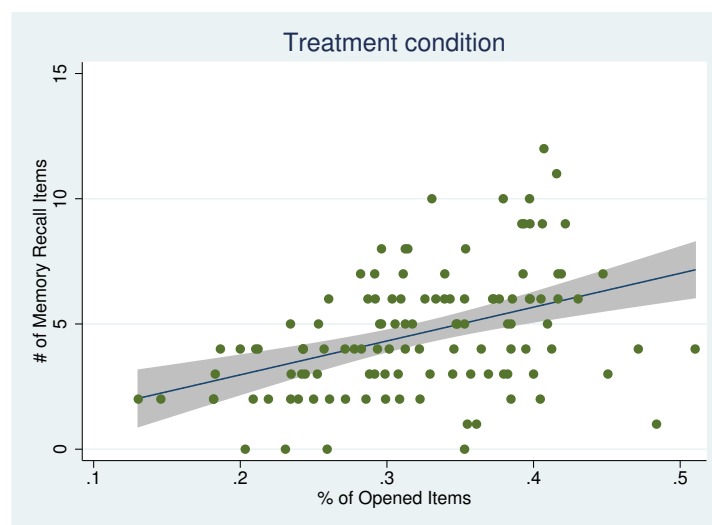
Figure 5.14: Marginal Effect of Treatment Across IAT Score, High Explicit Bias
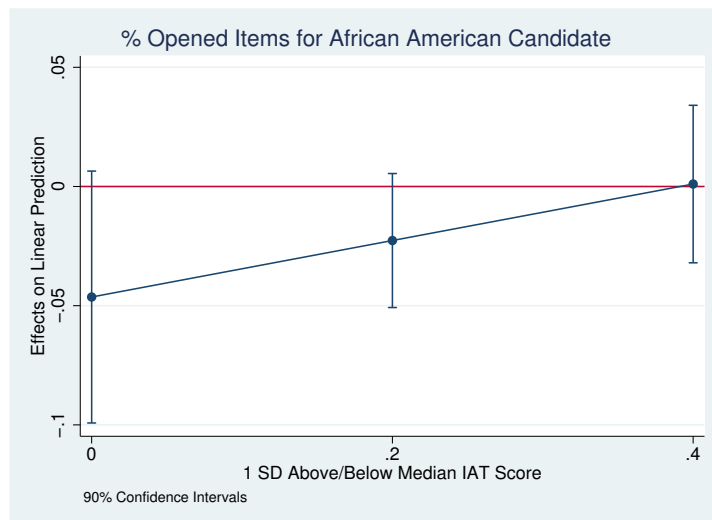
Figure 5.15: Marginal Effect of Treatment Across IAT Score, Low Explicit Bias
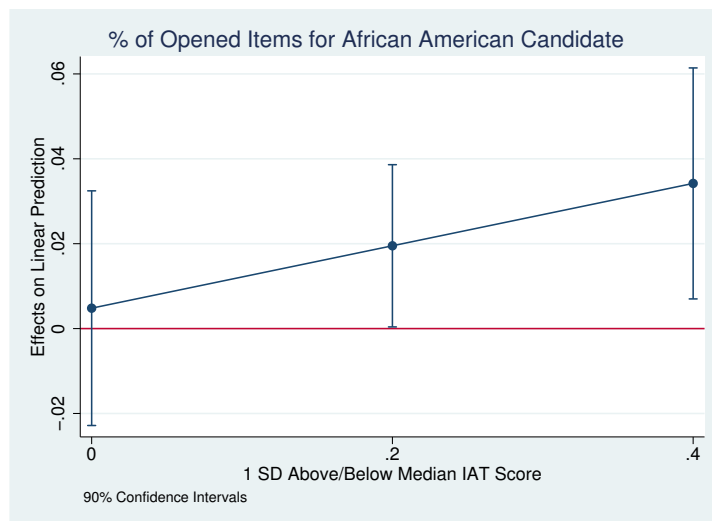
Chapter 6

Emotional state and the Activation of Implicit Racial Bias

6.1 Introduction

Previous research has shown that specific emotions can have an important effect on information processing and decision-making (Damasio, 1994, 1996; Lerner and Keltner, 2001; Lerner and Tiedens, 2001). In political science, a number of studies have demonstrated that different emotions have distinct effects on political information search (Brader, 2005; Brader et al., 2008; Huddy and Feldman, 2007; Lerner, 2001; Marcus et al., 2000; Parker and Isbell, 2010). Much of this work stems from Marcus and colleagues (1993, 2000) affective intelligence (AI) research program. AI make a distinction between two affective sub-systems of the brain: the disposition and surveillance system. The disposition system operates in familiar contexts and thus uses already established behavioral routines to deal with the external environment. The surveillance system, on the other hand, identifies stimuli that are unfamiliar and thus shifts attention away from pre-established routines. Using this framework, the AI research program has provided significant evidence demonstrating that specific emotions are more (and less) likely to trigger the activation of the disposition and surveillance system, respectively.

The goal of this paper is to incorporate the insights of affective intelligence into the study of implicit attitudes. There are several reasons to think that affective intelligence may help inform the literature on implicit attitudes and political decision-making. Perhaps most importantly, implicit attitudes are typically characterized as automatic affective evaluations. (Lodge and Taber, 2000, 2005, 2013; Perez, 2013). Consistent with this definition, neuroscientists have shown that implicit (but not explicit) attitudes are associated with activity in the amygdala, which is the part of the brain that handles

emotional responses to negative stimuli (Phelps et al. 2000; Cunningham, 2004). This close connection between implicit attitudes and the cognitive mechanisms associated with emotions raises several important testable hypotheses for how both of these factors influence political decision-making. If affective intelligence is correct to point out that positive and negative affect and specifically different sets of emotions are associated with different sub-systems of the brain (i.e. disposition vs. surveillance), then this raises questions as to whether affective reactions help to facilitate (or inhibit) the use of implicit attitudes in decision-making. If this indeed the case, then incorporating the insights from the affect and emotions literature into research on implicit attitudes may greatly improve our understanding of when such attitudes are likely to have their strongest (and weakest) effects on decision-making.

In this paper, I explore these issues using the dynamic information board methodology (Lau and Redlawsk, 1997; 2006) to construct a mock presidential primary campaign where subjects are presented a matrix of "learnable" information about fictional political candidates, which continuously flow down their computer screen for a limited length of time. This is followed by the administration of the "black-white" Implicit Association Test (IAT) as a measure of implicit racial bias. To test whether the impact of implicit attitudes on political information processing is moderated by emotional state, subjects were administered the Positive and Negative Affect Scale (PANAS-X) (Watson and Clark, 1999) one week prior to their arrival to the lab. The PANAS-X scale is a measure typically used to assess the two primary dimensions of emotional state or mood: positive and negative affect (Watson, Clark, and Telegen, 1988).[1]

In terms of results, I find, consistent with affective intelligence, that implicit attitudes have little impact on respondents' decision-making behavior for those self-reporting to be in a negative emotional state. Specifically, I find that the marginal effect of IAT (on information search) is indistinguishable from zero for those self-reporting

---

[1]For this study, I use the terms positive and negative affect in order to be consistent with the PANAS-X scale. The positive and negative affect scales draws from the same mood terms used by Marcus and Mackuen (2000) to construct their enthusiasm and anxiety scales. As such, for this study positive and negative affect and enthusiasm and anxiety are used interchangeably.

to be in a negative emotional state. In contrast, implicit attitudes have a strong statistically significant effect for those reporting to be in a positive emotional state. This is consistent with the hypothesis that a positive emotional state facilitates the use of implicit attitudes while a negative one helps to inhibit it.

Shifting to the relationship between explicit attitudes and emotional state, consistent with affective intelligence, I find that explicit attitudes are only a significant predictor of responder behavior for those reporting to be in a negative emotional state. Specifically, I find that the marginal effect for explicit prejudice is a statistically significant predictor of information search only for those in the negative emotional state. This is again consistent with the hypotheses derived from affective intelligence.

In sum, the findings from this study have several important implications for the growing literature on implicit attitudes and political decision-making. By demonstrating that the predictive strength of implicit attitudes varies with emotional state, my study helps to shed further light on the potential boundary conditions for understanding the influence of such attitudes on political behavior.

6.2 Literature Review

6.2.1 The role of affect in politics

A number of studies have demonstrated the impact of emotions and emotional state on public opinion and political behavior more generally (Marcus and Mackuen, 1993, 2000; Brader, 2005; Valentino et al, 2008; Redlawsk et al., 2007).[2] Much of this work stems from research in cognitive neuroscience and social psychology that documents two emotional sub-systems in the brain (the disposition and surveillance system), which

---

[2]Before moving forward with a discussion of this literature, it is important to make a distinction between emotions and emotional state (or mood). Psychologists typically distinguish emotions from emotional state by characterizing emotions as affective reactions to a stimuli or object whereas emotional state is generally thought of as affective reactions that are unfocused, diffuse, and usually temporally remote from it's cause (Batson, Shaw and Oleson, 1992; Feldman and Russell, 1999; Frijda, 2009). Given that this study records affect one week prior to the lab sessions and I am not measuring affect toward a particular stimuli, the focus of this study is on the impact of emotional state rather than emotions more generally.

trigger distinct cognitive mechanisms for dealing with exposure to new information. Specifically, the disposition system operates in familiar contexts and thus uses already established behavioral routines to deal with the external environment while the surveillance system, on the other hand, identifies stimuli that are unfamiliar and thus shifts attention away from pre-established routines. Turning to the role of affect in triggering the use of these different cognitive mechanisms, researchers have consistently found that emotional state (i.e. positive or negative affect) is systematically associated with the activation of distinct sub-systems of the brain. Negative affect, and specifically negative emotions such as fear and anxiety heighten threat levels thus facilitating the activation of the surveillance system, which results in increased levels of attention and information seeking for those individuals experiencing such emotions. On the other hand, positive affect and specifically positive emotions such as enthusiasm are associated with the triggering of the disposition system, which reinforces habitual behavior and thus is typically characterized by less active attention and information-seeking.

In political science, numerous studies have found that affective reactions to political candidates is associated with systematic differences in political information processing. Marcus and colleagues (1993, 2000) are largely credited with the first systematic study of affect and political behavior. In their studies, the authors rely on mass level panel survey data (ANES) to construct measures of emotional response to different political candidates. From these evaluations, they construct a two-dimensional measure of affect meant to capture positive and negative emotions. The authors go on to demonstrate a robust correlation between self-reported affective reactions to candidates and interest in politics, information search, and political learning. In line with their hypotheses, Marcus and colleagues report that self-reported negative affect about a candidate was associated with more active information search and greater knowledge about where candidates stand on a variety of issues (relative to those self-reporting positive affect). In sum, the authors conclude that negative affective reactions to candidates can actually lead individuals to pay closer attention to the political environment and enhance individuals abilities to learn more about the candidates.

Since Marcus and colleagues seminal work on affective intelligence, subsequent studies have tried to further unpack the relationship between emotions and political information processing. This in part stems from the fact that since Marcus et al's work relies on panel survey data, it is unclear whether affective reactions drive information search and political learning or whether it is the other way around (i.e. endogeneity bias). To account for this, recent studies have leveraged experimental designs that have sought to manipulate individuals' emotional state via random assignment. For instance, in Brader (2005) the author randomly assigns subjects to campaign ads cueing distinct emotions (i.e. enthusiasm and anxiety respectively) prior to participating in a mock election. In the enthusiasm group, subjects are exposed to an advertisement that emphasizes success and good times while in the anxiety group, the ad focuses on threatening imagery and includes black and white images depicting violence and drug use. As expected, Brader finds that subjects in the enthusiasm treatment expressed a greater interest in the campaign and exhibited a greater reliance on their pre-existing preferences to make their voting decision. Subjects in the anxiety group, on the other hand, engaged in more information-seeking behavior and were less likely to rely on their pre-existing beliefs.

Other studies have focused primarily on political information search. In both Valentino et al (2008) and Parker and Isbell (2010), the authors manipulate emotional state (via a common mood induction procedure) and explore whether the cueing of distinct emotions leads to systematic differences in information search patterns in the context of a mock election campaign conducted via an information board where subjects must search for information about each of the fictional candidates. In both studies, the authors report that inducing negative emotions such as fear and anxiety led to statistically significant increases in both the amount of information accessed and the time spent reading such information. In a separate study, Redlawsk et al (2007) conduct a mock election campaign using a dynamic information board and randomly assign emotional state by manipulating the congruency of information between the subject and their most preferred candidate. The idea being that the more incongruent subsequent

information is about a subject's most preferred candidate, the more anxiety that will be produced. Consistent with the other studies, the authors find (with some qualifications) that anxiety increases information processing time relative to their enthusiasm condition. Finally, Mackuen and colleagues (2010) created a web-based environment where subjects must search for information about an affirmative action policy. In this study, the authors find that anger induces more limited and selective information search, while anxiety leads to broader, more open information search.

6.2.2 Bringing in implicit attitudes

That a systematic investigation incorporating implicit attitudes in the study of emotion and politics has yet to be done is puzzling given the typical characterization of implicit attitudes as "automatic" affective associations (Lodge and Taber, 2000, 2005, 2013; Perez, 2013). In other words, implicit attitudes are generally thought to reflect instant valence judgements (e.g. do you associate the target with good/bad, favorable/unfavorable, etc) (Smith and Nosek, 2011; Payne et al., 2005; Spence and Townsend, 2008). Consistent with this definition, neuroscientists have shown that implicit (but not explicit) attitudes are associated with activity in the amygdala, which is the part of the brain that handles emotional responses to negative stimuli (Phelps et al. 2000; Cunningham, 2004). Hence, if we characterize implicit attitudes as automatic affective associations, the the question becomes whether the elicitation of affective responses can facilitate (or inhibit) the activation and use of implicit attitudes.

There is already some evidence to suggest that the activation of implicit attitudes can be facilitated by specific emotional states. In two studies, Destento and Dasgupta et al (Desteno et al., 2004; Dasgupta et al., 2009) incorporate emotions into the study of implicit attitudes. In Desteno et al (2004), the authors sort subjects into minimal in-group and out-groups. Subjects are then randomly assigned to a condition where one of three emotional states are induced: anger, sadness, or a neutral state. Finally, subjects are administered an IAT used to measure the strength of association between the in-group/out-group and positive/negative attributes. The authors hypothesize that

anger should facilitate automatic evaluations of the in-group/out-group since anger is generally associated with system 1 processing and a greater readiness to use cognitive heuristics (Bodenhausen, Kramer, Susser, 1994; Bodenhausen, Sheppard, Kramer, 1994; Tiedens Linton, 2001; Marcus and Mackuen, 2001). In line with their prediction, the authors report a significant interaction between their anger manipulation and IAT score. They found no such interaction with their other emotional cues. In a similar study, Dasgupta et al (2009) use the same experimental design but this time test it for known out-groups. Hence, in this study they use Arab-Americans as the relevant out-group and again find that the strength of implicit prejudice against Arabs increased when subjects were primed with the anger emotional cue.

This preliminary evidence provided by Desteno, Dasgpupta, and colleagues suggests that there may be substantial gains from incorporating the insights of affective intelligence into the study of implicit attitudes. Specifically, the findings from affective intelligence regarding when individuals are more (and less) likely to engage in deliberative thought (i.e. use of the surveillance system) may also help to inform us as to under what conditions implicit attitudes (and specifically implicit stereotypes) are likely to be used in decision-making. This will also help to deepen our understanding of some of the potential boundary conditions for implicit attitudes' impact on various forms of political behavior. Hence, the focus of this study will be to examine the relationship between emotional state and the activation (and use of) of implicit bias in the context of political information processing.

6.3   Research strategy

To examine the relationship between emotional state and implicit bias, I use the data on information processing from study 1 (see chapter 4) since this study contained no manipulations to the information environment and subjects in this condition were administered an emotional state survey instrument one week prior to their arrival to the laboratory. The experimental design is exactly the same as chapter 4. Upon arrival to the lab, subjects participate in a mock campaign and election administered via the

dynamic information board methodology (Lau and Redlawsk, 1997, 2006), which is then followed by the administration of the 'black-white' race IAT.

To measure emotional state, I use the PANAS-X scale (Watson and Clark, 1999). The PANAS-X scale has been used in a variety of different studies to explore the impact of affect, emotions, and mood on decision-making behavior (Friese et al., 2008; Lerner and Keltner, 2001; Redlawsk et al., 2010) and asks subjects to assess (using a five-point scale) words and phrases that describe different feelings and emotions. For instance, words such as afraid, scared, and nervous characterize a negative emotional state, while words like cheerful/excited/delighted characterize happiness/joviality describe a positive emotional state). For each word, subjects are asked to "indicate the extent you have felt this way during the past few weeks."[3] After subjects completed the PANAS-X survey measure, they scheduled a lab session for approximately one week later.[4]

Let's now spend a few minutes discussing the scale construction from the PANAS-X measure. To construct the emotional state measure, I use a modified version of the PANAS-X scale for positive and negative emotions.[5] As in the PANAS-X scoring procedure, I create summary scale items for positive and negative emotional state. Specifically, I sum the scores for each positive (and negative) PANAS-X survey item to create a additive scales for positive and negative affect. However, in order to better measure relative emotional state, I construct a difference score for each individual that

---

[2]More details on the experimental design can be read in chapter 4. Instead, I will focus attention in this section on the emotional state measure administered one week prior to the subjects' lab session.

[3]Please review the appendix in Chapter 3 for the exact wording of each survey item.

[4]Given the nature of the emotional state measure, there are several important differences between this study and previous ones exploring the relationship between affect and information processing. Unlike previous studies, the focus here is not on affective or emotional reactions to specific stimuli. Rather I administer the affect scale measure one week prior to the lab session as a way to measure emotional state. Recall that unlike emotions, emotional state is not directed toward a particular stimuli, moreover what I call emotional state or mood typically lasts much longer than emotions.

[5]Marcus and colleagues (2000, 2006) make the important point that emotion must be thought of as more than simple positive-negative valence. Instead, discrete emotions such as anxiety, enthusiasm, and anger may have differential effects on individual behavior. For this study, I chose to use a dichotomous positive-negative affective state scale in order to produce enough variance in the scale so as to recover effects. Since this study uses a relatively small sample, creating discrete emotion scales is difficult since they lack sufficient variance for recovering potential effects. However, as a robustness check, available upon request, all of the results of this study were re-run with removing all anger items from the negative affect scale. There are no substantive changes in the results with this alternative specification.

subtracts their negative affect score from their positive affect score. I do this to account for the fact that an individual may be above or below the median for both the positive and negative affect scale. As such, by not using a relative measure (i.e. taking a difference for the two scale scores), I gain little traction on emotional state.[6]

Hence, my modified emotional state scale is the difference in summary scale scores for positive and negative affect. The scale ranges from -34 to +27 with negative scores representing relatively high levels of negative affect and positive scores representing relatively high levels of positive affect. Figure 6.1 shows the distribution of scores for my affect scale. The scale is close to normally distributed with a mean very close to zero.

- Insert Figure 6.1 Here -

6.4   Hypotheses

Given the findings of affective intelligence theory and what I know about the relationship between implicit racial attitudes and political information processing (as demonstrated in Chapter 4), I can derive several testable hypotheses regarding the potential conditional effects of emotional state on the association between implicit racial bias and political information search.

H1: The effect of implicit prejudice on political information processing will be stronger for those self-reporting to be in a positive emotional state and weaker for those self-reporting to be in a negative emotional state.

This hypothesis is depicted visually in Figure 6.2 below.

----

[6]The logic behind using a relative measure of emotional state is similar to using a difference score for candidate feeling thermometers (Winter and Berinsky, 1999). Specifically, such measures are typically plagued by problems of inter-personal comparability. In other words, different individuals may interpret the scale in different ways, thus making it difficult to compare scores across individuals. As suggested by Winter and Berinsky (1999), one way to help account for this is to create an individual difference score between the scales of interest. Please note that alternative scoring strategies were also used and yield no substantively different findings for the main model specifications.

- Insert Figure 6.2 Here -

For implicit prejudice, I expect that the marginal effect for IAT should be close to zero for those in a negative emotional state. Since negative emotions are associated with more deliberative processing (i.e. more information search), I should see IAT's effect be relatively muted for those in a negative emotional state. Conversely, for those in a positive emotional state, I expect more of a reliance on cognitive heuristics and the disposition system. For these individuals, the marginal effect of IAT should be much stronger as compared to those possessing negative emotions.

H2: The effect of explicit prejudice on political information processing will be stronger for those self-reporting to be in a negative emotional state and weaker for those self-reporting a positive emotional state.

For explicit prejudice, I expect that the marginal effect for explicit prejudice[7] to be larger for those in a negative emotional state since they are presumably relying more on their surveillance system, thus using greater deliberative thought. On the other hand, I expect the effect of explicit prejudice to be weaker with individuals in a more positive emotional state since these individuals are likely relying more on their disposition system, thus allowing implicit mechanisms to guide information processing behavior. This hypothesis is depicted visually in Figure 6.3 below.

- Insert Figure 6.3 Here -

H3: The relationship between emotional state and information search will be conditional on the strength of the implicit attitude.

For this final hypothesis, I am interested in examining whether implicit attitudes can help to moderate the relationship between emotional state and information search. In

---

[7]In this study, I report results for explicit prejudice using the symbolic racism scale. However, all models were re-run with the ethnocentrism scale as a robustness check. I find no substantive differences in results from these models. Results from these models can be found in this chapter's data appendix.

other words, it may be case that the effect of emotional state on information processing can be overriden by those with high levels of implicit bias. In this case, negative emotional state may have little effect on political decision-making for those with high levels of implicit racial bias.

6.5   Results

Our examination of emotional state, implicit prejudice, and information processing focuses first on the relationship between emotional state and information search. As mentioned earlier, a number of studies report that negative emotions or affect tend to be associated with the activation of the surveillance system and thus is associated with higher levels of information-seeking. Conversely, positive emotional state is associated with the disposition system and as such is typically associated with less active information search.

I can further examine the robustness of these claims by analyzing the relationship between the total number of opened items (across all candidates) and the emotional state scale. Figure 6.4 shows a density plot for the total number of opened items for those above and below the median emotional state score. Consistent with affective intelligence theory, I see that those below the median emotional state score (i.e. negative emotional state) searching for more information compared to those above the median (i.e. positive emotional state). This difference is very close to reaching statistical significance.[8]

- Insert Figure 6.4 Here -

Another way to explore the main effect of emotional state on information processing is to model the total number of opened items as a function of emotional state and our other predictors associated with information search. Table 6.1 provides the results for

---

[8]I conducted the same set of analyses for total time duration but I fail to recover any significant differences.

a series of regressions where I do just that. In model 1, I use as an outcome variable the total number of opened items across all candidates. In models 2 through 4, I use the total number of opened items for each candidate.[9]

- Insert Table 6.1 Here -

Interestingly, I find that while the emotional state measure moves in the expected direction (i.e. consistent with affective intelligence theory) in each of my models, it is only a significant predictor of information search in the model for the African American candidate. Given data limitations, I am generally unable to unpack this finding. However, one potential hypothesis relates to recent findings (Redlawsk et al., 2007; Valentino et al, 2008), which report that the effect of emotions will likely impact information processing for the candidate that is the source of emotional response. If this is indeed the case, then my findings may suggest that the African American candidate is helping to cue an emotional response (above and beyond my pre-recorded measure of subjects' emotional state), however more research is needed to explore such an hypothesis.[10]

Let's now turn our attention to whether I can identify a meaningful relationship between implicit prejudice and emotional state. Specifically, let's first examine whether the relationship between implicit (and explicit) prejudice and information search for the African American candidate is moderated by emotional state. Recall that my working hypothesis is that negative affect should help to activate the surveillance system, thus inducing more active information search and consequently crowding out the potential effects of implicit prejudice and perhaps amplifying more deliberative forms (i.e. explicit) of prejudice. Relatedly, I expect that a more positive emotional state should be associated with the disposition system, thus inhibiting more deliberative forms of information search and instead facilitating the use of implicit attitudes.

---

[9]Please note that I am focused here on depth of information search rather than comparability. For that reason, the dependent variable used here is the raw number of opened items rather than a percentage of total opened items.

[10]Please see the appendix, table A1 for these same set of regressions for my time duration variable.

Table 6.2 provides my primary model specification where I regress the two informa-tion processing variables (i.e. total opened and total duration) for the African American candidate on the interaction between the implicit and explicit prejudice items and the emotional state measure, while controlling for the standard set of covariates. This model specification yields several interesting results.

- Insert Table 6.2 Here -

For both models, I see that the main effect of both prejudice items (implicit and explicit) along with the emotional state measure are all significant and in the expected direction. However, in order to understand the nature of the interaction effects, I must plot these results. A first way to examine potential conditional effects is to simply compare the slope of IAT (i.e. IAT's effect on the dependent variable) for those self-reporting to be in a positive and negative emotional state. If the interaction between emotional state and implicit/explicit prejudice operates as expected, then I should see differences in IAT's slope as a function of emotional state.

Figure 6.5 presents the IAT effect for those reporting to be in positive and negative emotional state respectively. Specifically, I examine the slope of IAT when emotional state is one SD below the mean (i.e. negative affect) and one SD above the mean (positive affect) respectively.

- Insert Figure 6.5 Here -

I see that while the slopes are very similar, they differ in the location of the intercepts. Specifically, for negative affect the intercept is greater (relative to those with positive affect), thus suggesting that IAT's effect is more muted for those in a negative emotional state. Turning to slopes for explicit prejudice (figure 6.6), I observe results very much consistent with my hypothesis.

- Insert Figure 6.6 Here -

I see that explicit prejudice's effect for negative and positive affect are quite different. Specifically, for positive affect I see a flat line for explicit prejudice indicating essentially zero effect. In contrast, for negative affect, I find explicit prejudice to have a strong negative effect. In sum, these initial results suggest a potentially important conditional relationship between emotional state and the use of implicit and explicit prejudice.[11]

To further examine these results, I can plot the regression results in a slightly different way. Rather than just comparing slopes, I can instead plot the marginal effect of IAT across the emotional state scale. This will allow us to examine whether the effect of implicit bias on information processing varies as a function of emotional state. Figure 6.7 presents the results for the marginal effects plot.

- Insert Figure 6.7 Here -

Consistent with my hypotheses, I find that the marginal effect of IAT is indistinguishable from zero at low levels of the emotional state measure (i.e. negative affect), however this changes once I move toward a more positive emotional state. I see that as I move toward higher positive affect, the marginal effect of IAT strengthens and is distinguishable from zero across a number of different (positive) values. This is in line with the hypothesis that positive emotional state facilitates the use of implicit attitudes while negative affect acts to inhibit the use of such attitudes.

Finally, I must also consider whether the impact of explicit prejudice or stereotypes also vary depending on emotional state. Related to the discussion above regarding the conditional nature of implicit prejudice, I also hypothesized that the use (and thus the predictive strength) of explicit prejudice should be greater when individuals are induced to rely more on their surveillance system (i.e. when negative affect is activated)

---

[11]See the appendix, figures A6.1 and A6.2 for these same plots with the time duration variable.

and weaker when positive affect is induced (i.e. when positive affect is activated). Figure 8 displays the same marginal effects plot but now for the explicit prejudice measure. I see that the marginal effect for explicit prejudice is distinguishable from zero only at very low scores on the emotional state scale (i.e. negative affect). Again, this is consistent with the hypothesis that explicit prejudice is a stronger predictor of information processing behavior when individuals' surveillance system is activated via negative affect.[12]

- Insert Figure 6.8 Here -

For my final set of analyses, I turn my attention to examining whether the effect of emotional state is conditional on the strength of implicit attitudes. In other words, are those with highly crystallized implicit attitudes better able to override the effects of emotional state on decision-making behavior? To test this hypothesis, I can use the same baseline regression model where I regress the percentage of opened items for the African American candidate on the standard set of covariates but now I can include an interaction between my dichotomized IAT measure (i.e. a "1" for those above the median IAT score and a "0" for those below the median IAT score) and the continuous emotional state scale measure. I can then plot the predicted values of the outcome variable across the emotional state scale for low and high IAT individuals respectively. Table 6.3 provides the regression results and figure 6.9 plots the predicted values for the interaction.

- Insert Table 6.3 and Figure 6.9 Here -

If emotional state operates the same way for low and high implicit bias individuals, then I should see that in the plotting of the predicted values, a complete overlap for low and

---

[12]I also ran the same set of models for the second information processing variable, total time spent reading information about the African American candidate. I find qualitatively the same results, however the reported conditional effects are weaker. That being said, all of the models are in the expected direction and yield results quite similar to the total opened variable. Results for these models can be found in the appendix, figures A6.3 and A6.4. Please note that these results are qualitatively similar to that of the total opened variable.

high IAT subjects. However, this is not what is observed in figure 6.9. Instead, I find that across most values of the emotional state scale, high implicit bias individuals search for less information about the African American candidate. This operates at both low end of the scale (i.e. negative affect) and the high end of the scale (i.e. positive affect). Such results suggest that emotional state has little impact on information processing for those with high levels of implicit bias. Rather, most of the effect of emotional state operates for those with relatively low levels of such bias.

Table 3 also provides the same regression for the total reading time for the African American candidate (model 2) and Figure 6.10 displays the predicted values plot. Again, I observe the same pattern of results with high IAT individuals consistently spending less time reading information about the African American candidate at all levels of the emotional state scale.

- Insert Figure 6.10 Here -

In sum, these results suggest that while emotional state can have an important moderating effect on implicit bias, the strength or crystallization of those implicit attitudes make a difference. That those with high levels of implicit bias are largely able to override the effect of emotional state suggests that there are limits to the influence of affect on implicit attitudes.

6.6   Conclusion

As the literature on implicit biases and political behavior continues to grow, more attention will be paid to the pathways through which such attitudes influence the decision-making process. A number of studies have already proclaimed that the next frontier for implicit attitude research in political science (and more broadly) is to examine potential moderators of implicit processing in an attempt to understand the conditions under which these attitudes are likely to have their strongest (and weakest) effects on individual behavior. The focus of this paper was to examine one such

moderator in the form of emotional state. An already robust body of literature in political science has demonstrated the strong relationship between emotions and political information processing (Marcus and Mackuen, 1993, 2000; Brader, 2005; Redlawsk et al., 2007; Valentino, 2008). However, to date there has been little investigation into whether the incorporation of emotional state into the study of implicit attitudes yields unique insights in understanding when implicit attitudes are likely to be activated and subsequently influence behavior.

In this paper, I used an experimental design where I recorded subjects' emotional state (positive or negative) one week prior to their arrival to the lab. Once at the lab, subjects participated in a mock presidential primary campaign and election using the dynamic process tracing methodology (Lau and Redlawsk, 1997, 2006). The results from this study provide several important implications for the study of implicit attitudes. Consistent with the literature from affective intelligence, I find that the impact of implicit bias on political information processing does appear to vary by emotional state. Specifically, my results show that the effect of implicit racial bias on information processing was much stronger for those self-reporting as being in a positive emotional state. In fact, I fail to recover a significant IAT effect for any of the information processing variables for those individuals self-reporting as being in a negative emotional state. In contrast, I find that more explicit (and deliberative) forms of racial bias operate in the opposite way with the predictive strength of explicit racial attitudes having a much larger effect on information processing for those in the negative affect group (compared to the positive affect group). Finally, I also report that the main effect of emotional state is itself conditional on the strength of one's implicit bias. In other words, I provide some evidence to suggest that the effect of emotions on information search can be overridden if one's implicit bias is sufficiently strong.

In sum, the results from this study shed further light on the potential pathways through implicit attitudes influence behavior. By highlighting the conditional relationship between emotional state and the predictive strength of implicit bias, this study provides a better understanding as to when implicit bias is likely to be a significant

driver of behavior. With regards to the implications for future research in political science, my findings suggest that affective responses to political stimuli may play an important role in the activation and use of implicit attitudes. Such findings also suggest that research on the influence of media effects and political communication may be an increasingly important domain for understanding the influence of implicit bias on political decision-making. To the extent that political communications can successfully activate seemingly discrete emotional states (positive vs. negative), the impact of implicit attitudes on political decision-making may very well vary. Thus, one goal of future research on implicit attitudes and political behavior should be to further unpack the relationship between emotional state and implicit attitude activation.

6.7   Tables and Figures

Table 6.1: The impact of emotional state on total number of opened items

	All	African American	White Moderate	White Partisan
IAT	3.190	-3.842✝	-2.805	8.525**
	(15.966)	(3.077)	(3.353)	(4.422)
Explicit Bias	4.331✝	-0.184	1.175*	-0.537
	(3.008)	(0.623)	(0.671)	(0.802)
Emotional state	-0.132	-0.072✝	0.045	-0.022
	(0.308)	(0.055)	(0.045)	(0.074)
Controls	Yes	Yes	Yes	Yes
Constant	75.133***	-100.953	-93.942	-97.690
	(15.416)	(10.851)	(10.221)	(11.273)
N	115	115	115	115
R^2	0.041	0.707	0.679	0.588

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]All models were run using robust standard errors. Controls for models include gender, race, political knowledge, partisanship, the total number of opened items (only included in the candidate specific models), the total reading time, and reading speed

Table 6.2: Information processing for African American candidate, across emotional state

	Total Opened	Total Duration
IAT	-0.061*	-0.072✝
	(0.035)	(0.049)
Explicit Bias	-0.010✝	-0.018*
	(0.006)	(0.009)
Emotional state	-0.003*	-0.005**
	(0.001)	(0.002)
Emotional state * IAT	-0.002	-0.001
	(0.003)	(0.004)
Emotional state*Explicit	0.001**	0.001**
	(0.0005)	(0.0007)
Controls	Yes	Yes
Constant	0.259***	0.183
	(0.088)	(0.177)
N	115	115
R^2	0.150	0.134

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]All models were run using robust standard errors. Controls for models include gender, race, political knowledge, partisanship, the total number of opened items, the total reading time, and reading speed

Table 6.3: Information processing for African American candidate, by IAT

	Total Opened	Total Duration
Above Median IAT	-0.032**	-0.043**
	(0.012)	(0.049)
Above Median Explicit Bias	-0.008	-0.023
	(0.013)	(0.018)
Emotional state	-0.001†	-0.003**
	(0.001)	(0.001)
Above Median IAT * Emotional state	0.001†	0.003**
	(0.001)	(0.001)
Above Median Explicit Bias * Emotional state	0.0002	0.001
	(0.001)	(0.001)
Controls	Yes	Yes
Constant	0.254***	0.110
	(0.090)	(0.168)
N	115	115
R^2	0.140	0.150

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]All models were run using robust standard errors. Controls for models include gender, race, political knowledge, partisanship, the total number of opened items, the total reading time, and reading speed

Figure 6.1: Density plot, emotional state scale
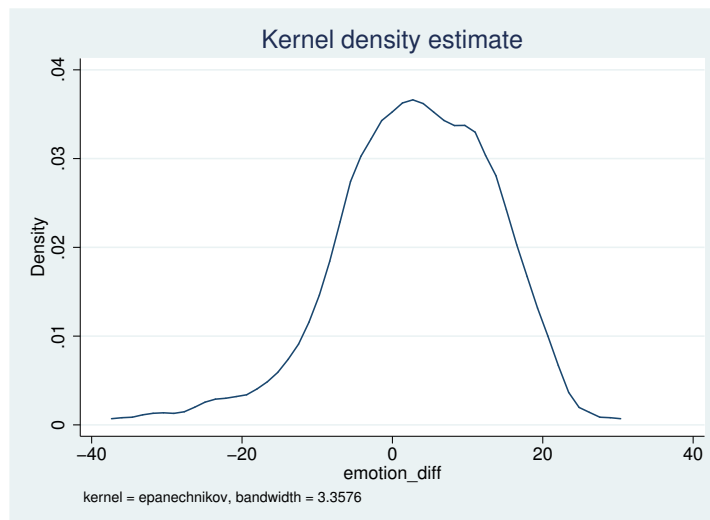
Figure 6.2: Hypothesized conditional effect of emotional state on implicit prejudice
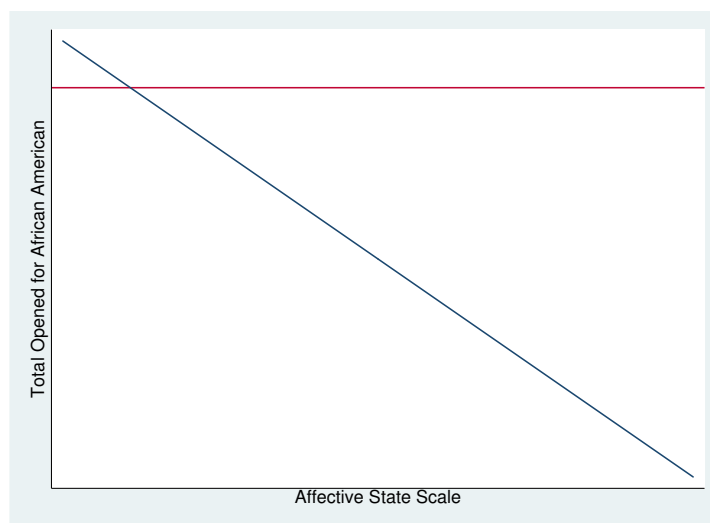
Figure 6.3: Hypothesized conditional effect of emotional state on explicit prejudice
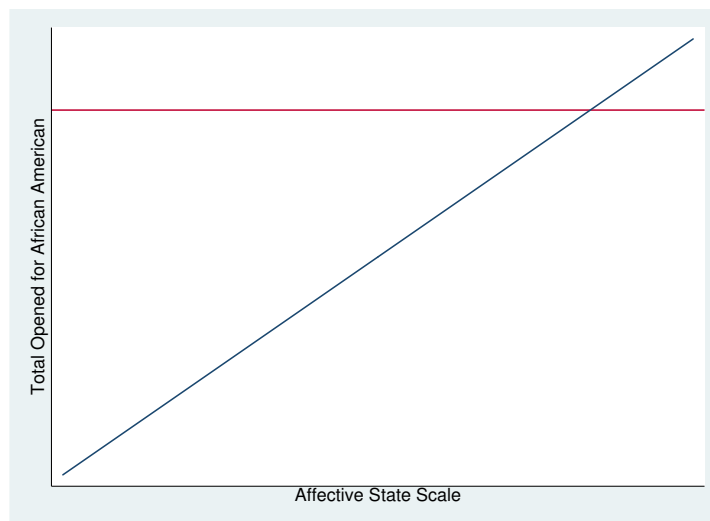
[a]Please note that the red line in both figures represents a zero effect.

Figure 6.4: Density plot for total number of opened items across emotional state scale
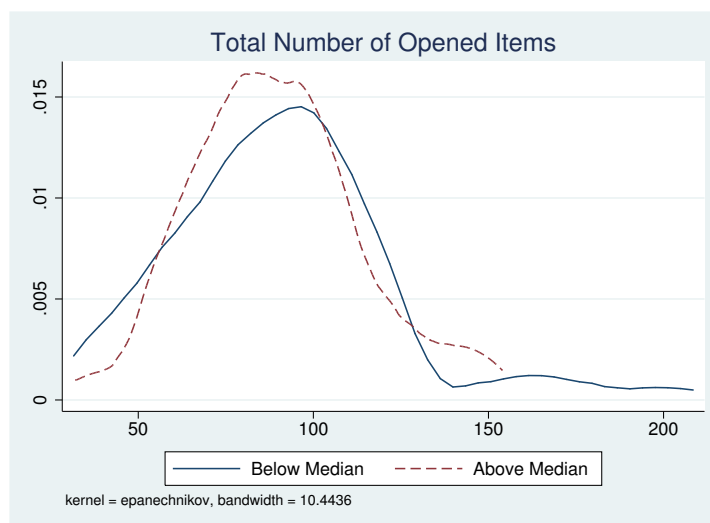
Figure 6.5: IAT effect across emotional state scale
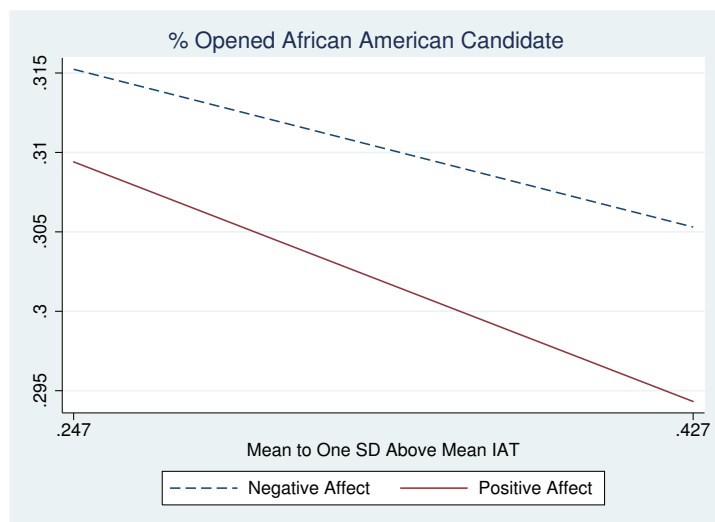
Figure 6.6: Symbolic racism effect across emotional state scale
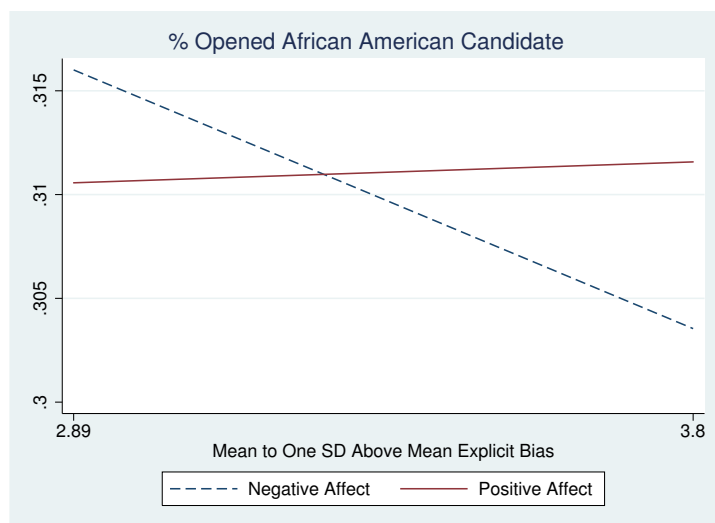
Figure 6.7: Marginal effect of IAT across emotional state scale

% Opened African American Candidate

90% Confidence Intervals
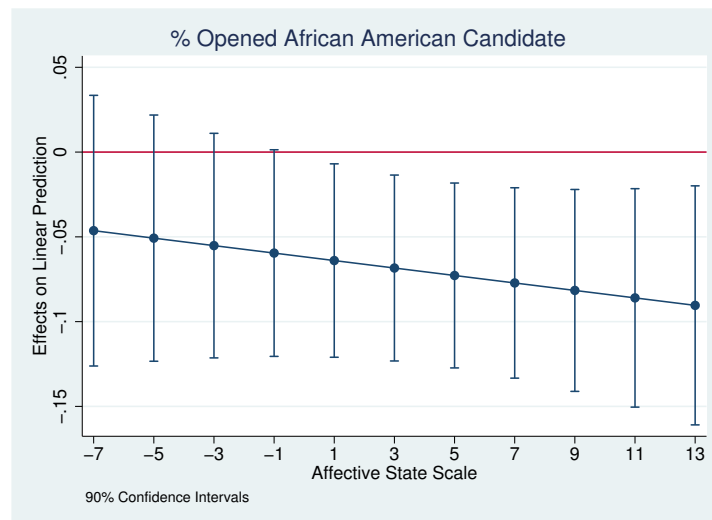
Figure 6.8: Marginal effect of Symbolic Racism across emotional state scale

% Opened African American Candidate

90% Confidence Intervals
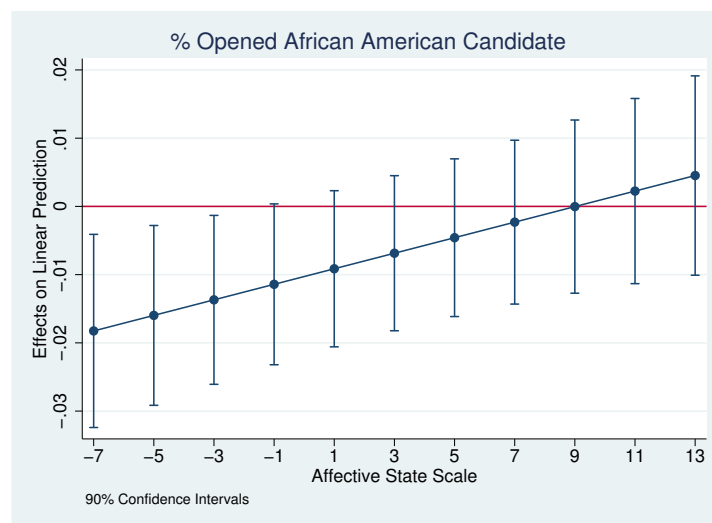
Figure 6.9: IAT effect at different levels of emotional state, Total opened African American candidate
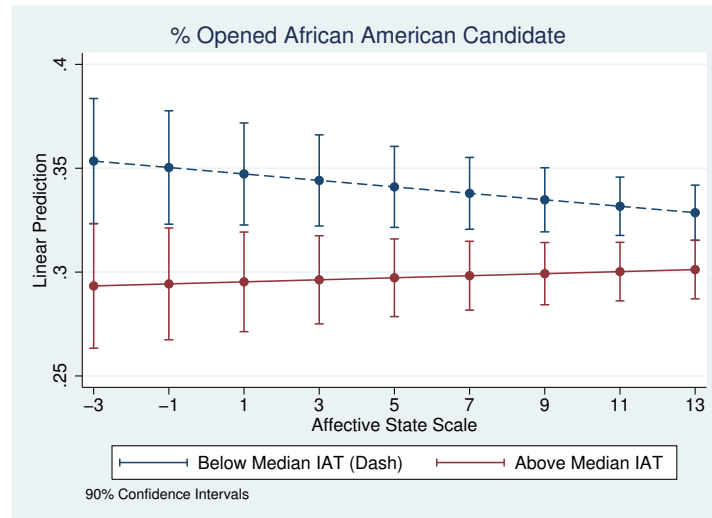
Figure 6.10: IAT effect at different levels of emotional state, Total duration African American candidate
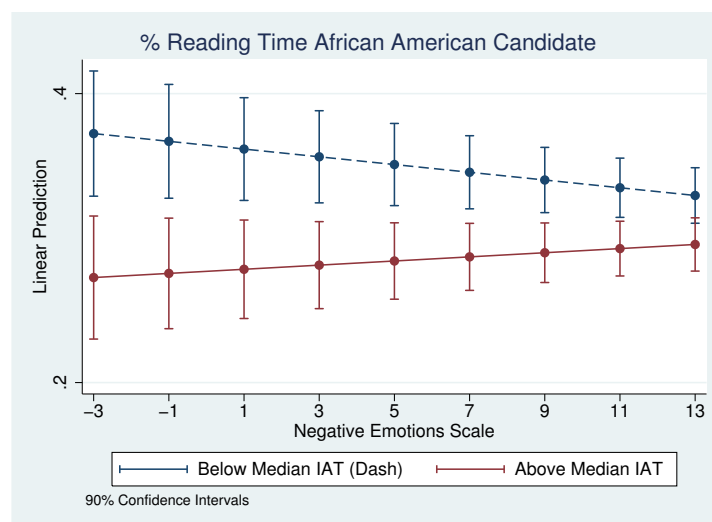
Chapter 7
Conclusion

The goal of this dissertation was to unpack the relationship between implicit attitudes and voter decision-making. A growing body of literature in political science has shown implicit attitudes to be a strong predictor of a variety of political outcomes. However, much less attention has been paid to the pathways through which such attitudes influence political decision-making. This has led to a recent debate in both psychology and political science on whether implicit attitudes lose predictive power once we account for more explicit forms of attitudes and preferences. Hence, the primary motivation for this project was to shed further light on the mechanisms linking implicit attitudes to political decision-making. In this final chapter, I will provide a brief overview of the objectives laid out in this dissertation project and discuss how the findings from each of my studies help to address these objectives. This chapter will conclude with a summary of the implications from this project and potential areas of future research.

7.1 Objective 1: Do implicit attitudes impact political information search?

The first objective of this dissertation was to design a study capable of directly assessing the relationship between implicit attitudes and political information search. Previous studies examining the relationship between implicit attitudes and information search have either failed to include a direct measure of implicit attitudes or have used a crude measure of information search that fails to account for a number of dimensions (see chapter 2 for an overview). To address these gaps, this dissertation provided a novel experimental design that used the Implicit Association Test (IAT) in conjunction with Lau and Redlawsk's (1997, 2006) dynamic process tracing methodology.

Using this research design, study 1 of this dissertation project focused on examining whether implicit racial bias impacted information search patterns for political candidates. Specifically, subjects were asked to participate in a mock Presidential primary election that included two white and one African-American fictional candidates. During the campaign, subjects searched for information (on a variety of different dimensions) about the candidates and at the conclusion, they are asked to evaluate each candidate and make a decision as to who to vote for. Once the mock campaign and election was complete, all subjects were administered the black-white race IAT to be used as a measure of implicit racial bias against African Americans.

The results from study 1 provide consistent evidence to support the claim that implicit attitudes help drive selection information processing. Implicit racial bias was shown to have a robust negative effect on information search for the African American political candidate. Specifically, implicit racial bias had a negative effect on both the amount of opened information and reading time duration for the African American candidate. I reported no such effect for the two white candidates in the primary campaign. A second set of important results from study 1 focused on the pathways through which implicit attitudes impact voter decision-making behavior. Similar to other studies, I found implicit racial bias to be negatively associated with both feeling thermometer scores and vote choice for the African American candidate. However, study 1 demonstrated that this negative effect loses significance once I account for information search patterns. These findings suggest that implicit bias may operate primarily as indirect effect on voter decision-making via biased information search.

7.2 Objective 2: Do manipulations of the information environment condition the effect of implicit attitudes?

The second objective of this dissertation project focused on examining potential moderators of implicit attitudes. Previous studies have largely failed to take seriously the potential *conditional* nature of implicit attitudes' effect on political behavior. Instead,

this literature has focused predominantly on demonstrating that such attitudes impact political decision-making. This is in contrast to a large body of literature in cognitive psychology focusing on the dual nature of information processing and the conditions under which these different systems are activated. Dual process theorists characterize information processing as being divided into two related but distinct systems: system 1 (implicit) and system 2 (explicit) processing. Furthermore, they contend that system 1 (or implicit) processing will have a larger impact on behavior in situations where time and resources to consciously deliberate are limited.

To empirically assess whether the insights from dual process theory can help inform the literature on implicit attitudes and political decision-making, study 2 of this dissertation project used the same experimental design as study 1, but this time introduced a manipulation that sought to vary a key feature of the information environment: cognitive capacity. In this study, a subset of subjects were randomly assigned to a condition where they were tasked with completing a number memorization task while searching for information about each of the presidential candidates.

The results from study 2 provide some evidence in support of dual process theory. Increased attention and engagement with the decision task helped to crowd out the influence of implicit attitudes on both political information search and candidate evaluation/vote choice. Specifically, implicit racial bias was shown to be a statistically significant predictor of both political information search and candidate evaluation only for those that paid less attention and were less engaged with the decision task. Conversely, for explicit prejudice I observed the exact opposite pattern. I found statistically significant effects only for individuals with high levels of attention and motivation. In total, these findings are consistent with the conjectures from dual process theory that implicit attitudes (or system 1 processing) are better predictors of behavior when individuals are less motivated to engage in deliberative thought.

7.3 Objective 3: Does emotional state moderate the effect of implicit attitudes?

For the third objective of this dissertation project, I shift focus from informational constraints to potential individual level moderators of implicit attitudes. Previous research has shown that specific emotional states can have an important effect on information processing and decision-making. Negative emotions (such as fear/anxiety/hostility) have been shown to trigger more active information processing and a shift away from easily accessible cognitive routines. Conversely, positive emotions have been shown to lead to increased heuristic use and more reliance on standard routines and stereotypes. Importantly, these insights have been largely ignored in the implicit attitudes literature with little empirical assessment of whether emotional states help to moderate or facilitate the activation of implicit attitudes.

To examine whether distinct emotional states help to moderate the impact of implicit attitudes, study 3 employed the same experimental design (as in study 1 and 2) but also included a measure of emotional state by examining subjects' score on the Positive and Negative Affect Scale (PANAS-X), which was administered one week prior to subjects' lab sessions. The PANAS-X scale consists of twenty emotion items on which participants indicate their present feelings (1=very slightly or not at all, 5=extremely). Summary scales were created to assess positive and negative affective states.

Consistent with previous studies on the relationship between emotions and information processing, the results from study 3 showed that the impact of implicit bias on political information processing does appear to vary by emotional state. Specifically, my results show that the effect of implicit racial bias on information processing was larger for those self-reporting as being in a positive emotional state. In fact, I recover a significant IAT effect for information processing for only those individuals self-reporting to be in a positive emotional state. For those self-reporting to be in a negative emotional state, the IAT effect is indistinguishable from zero. In contrast, more explicit (and deliberative) forms of racial bias operate in the opposite way with the predictive strength of explicit racial attitudes having a larger effect on information processing for

those in the negative affect group (compared to the positive affect group).

7.4   Implications

The findings from this dissertation project yield a number of important contributions to the burgeoning literature on implicit attitudes and political decision-making.

First, my findings have important implications for much of the work in political psychology and public opinion on the mechanisms linking implicit attitudes to voter decision-making. Unlike previous studies, this dissertation project offers an experimental design capable of directly examining the impact of implicit attitudes on political information search and vote choice. By providing important insight into the relationship between implicit attitudes and information processing, this study unpacks how automatic associations enter the decision stream. Moreover, by having individual level data on implicit attitudes, information search, and vote choice, this study is uniquely positioned to address how implicit attitudes can indirectly impact voter decision-making through biased information search.

The second contribution made by this dissertation involves an in-depth examination of some potential moderators of implicit attitudes impact on voter decision-making. By varying cognitive load and examining emotional state (respectively), this experimental design provides an empirical assessment of dual process theory and offers important insight on the potential boundary conditions for understanding the relationship between implicit attitudes and political behavior. Additionally, my finding that both cognitive capacity and emotional state impact the predictive strength of implicit attitudes represents an important area of future research for political scientists. Perhaps most importantly, these findings suggest that researchers must consider the potential differential effects of implicit attitudes based on features of the information environment.

Finally, the results from this project help inform the growing debate on the significance of implicit attitudes as predictors of political behavior. While there has been recent debate over whether implicit attitudes provide additional explanatory power to

models of voting behavior, the findings from this study suggest that such questions may be difficult to answer given that implicit attitudes likely enter the decision stream far earlier than vote choice. The findings from this study suggest that implicit attitudes impact more subtle forms of political behavior, specifically how individuals search for information about politics. Such findings are consistent with previous studies that have found that implicit attitudes tend to have stronger effects on non-verbal forms of behavior. As such, focusing attention only on the direct effects of implicit attitudes on vote choice may be misplaced.

7.5   Areas of future research

There are a number of important areas of future work on implicit attitudes and political behavior that can be explored as extensions to this dissertation project. Perhaps most importantly, future work is needed to understand the potential normative implications of the important role that implicit attitudes play in voter decision-making. The findings from this dissertation project have, in part, demonstrated that voters process information and form judgments subconsciously, which subsequently goes on to help structure explicit forms of behavior such as information search and vote choice.

Such conclusions raise important questions for the quality of democracy. From a normative standpoint, the fact that implicit attitudes help predict information search patterns and electoral success is troubling, particularly for those that view voter decision-making as guided by significant deliberation based on the gathering of numerous sources of information on each candidate. In fact, researchers have already started to question whether seemingly snap-judgments based on automatic affective associations are healthy for a functioning democracy (Olivola and Todorov, 2010; Riggio and Riggio, 2010). If voters are not taking the time to learn about all of the candidates and their issue positions due to their automatic or implicit attitudes, then this may be reflected in the quality of their vote choice.

However, while these potential implications are troublesome, it still remains an open

question as to how implicit attitudes impact the *quality* of voter decision-making. In fact, there is some research suggesting that such "gut feelings" may actually improve the quality of decision-making in certain contexts. A growing literature in social psychology suggests that it is not always beneficial to engage in intense conscious deliberation before making a decision. Contrary to popular belief, numerous studies have found that in complex decision environments, unconscious thought actually produces higher quality choices relative to deliberative decision-making (Dijksterhuis, 2004b; Dijksterhuis, Bos, Nordgren, van Baaren, 2006; Dijksterhuis Meurs, 2006; Dijksterhuis van Olden, in press; Nordgren Dijksterhuis, 2006).

Such findings raise questions of whether intense deliberation actually increases the likelihood that an individual makes a sub-optimal vote choice? In other words, should voters trust their "gut" more? Hence, the goal of future research on implicit attitudes and voter decision-making should focus more on the potential normative implications of using automatic attitudes to inform decision-making. How implicit attitudes impact the quality of vote choice is an empirical question that should be further researched. This differs from previous studies on the relationship between "automatic" processing or trait-based inference and political behavior in that the focus is not on demonstrating that such a link exists, but rather whether such rapid processing actually impacts the quality of voter decision-making.

Chapter 8

Appendicies

8.1 Appendix for Chapter 3

IAT Instructions (Monetary Incentive)

This part of the study will take about 15 minutes. The following is a list of category labels and the items. You will be presented with a set of words or images to classify into groups.

This task requires that you classify items as quickly as you can while making as few mistakes as possible. Going too slow or making too many mistakes will result in an uninterpretable score.

The following is a list of category labels and the items.

Category	Items
Good	Joy; Love; Peace; Wonderful; Pleasure; Glorious; Laughter
Bad	Agony; Terrible; Horrible; Nasty; Evil; Awful; Failure
African American	Faces of African American People
European American	Faces of European American People

KEEP IN MIND

- Keep your index fingers on the "e" and "i" computer keys to enable rapid response.
- Two labels at the top will tell you which words or images go with each key.
- Each word or image has a correct classification. Most of these are easy.
- The test gives no results if you go slow – Please try to go as fast as possible.
- Expect to make a few mistakes because of going fast. That's OK.

MONETARY REWARD
There is the chance to earn a monetary award for performing this task. We will be assessing how quickly and accurately you are able to sort items into categories. Those with sorting times in the top 20% will be awarded an additional $12 (this is in addition to your $8 show-up fee).

PLEASE NOTE: Sorting times are calculated using two components: (1) how quickly

you are able to classify items into categories and (2) the total number of correctly sorted pairs that you make.

Our scoring system to determine the top 20% will reward you based on two components: speed and accuracy.

Please keep in mind that our scoring system will penalize you if you focus on only one component (either speed or accuracy) exclusively.

In order to have a chance at the monetary reward, you must try to maximize both speed and accuracy equally.

Cognitive load condition

Throughout the primary, you will receive announcements. In each of these announcements, you will receive a five-digit number. Your task is to memorize each of these five-digit numbers.

Approximately seven minutes after each of these announcements, you will be asked to recall the five-digit number that you received.

Please note that you are not permitted to write down these numbers. You must try to memorize them.

The first random number sequence for you to memorize is: XXXXX

PANAS-X Scale

This scale consists of a number of words and phrases that describe different feelings and emotions. Read each item and then mark the appropriate answer in the space next to that word. Indicate to what extent you have felt this way during the past few weeks. Use the following scale to record your answers:

(1) very slightly or not at all; (2) a little; (3) moderately; (4) quite a bit; (5) extremely

1. afraid ( )
2. nervous ( )
3. angry ( )
4. scornful ( )
5. happy ( )
6. cheerful ( )
7. lively ( )
8. scared ( )
9. jittery ( )
10. hostile ( )
11. disgusted ( )
12. joyful ( )
13. excited ( )
14. energetic ( )
15. frightened ( )
16. shaky ( )
17. irritable ( )
18. loathing ( )
19. delighted ( )
20. enthusiastic ( )

Symbolic racism items

The items used to construct the symbolic racism scale are as follow:

1. Irish, Italians, Jewish, and many other minorities overcame prejudice and worked their way up. Blacks should do the same without any special favors.

2. Over the past few years blacks have gotten less than they deserve.

3. Its really a matter of some people not trying hard enough; if blacks would only try harder they could be just as well off as whites.

4. Generations of slavery and discrimination have created conditions that make it difficult for blacks to work their way out of the lower class.

Ethnocentrism scale

Now I have some questions about different groups in our society. Im going to show you a seven-point scale on which the characteristics of people in a group can be rated.[1] In the first statement a score of 1 means that you think almost all of the people in that group are "hard-working." A score of 7 means that you think almost all of the people in the group are "lazy. A score of 4 means that you think the group is not towards one end or the other, and of course you may choose any number in between that comes closest to where you think people in the group stand.

1. Where would you rate whites in general on this scale?

2. Where would you rate blacks in general on this scale?

3. Where would you rate Asian Americans in general on this scale?

4. Where would you rate Hispanic Americans in general on this scale?

Im going to show you a second a seven-point scale on which the characteristics of people in a group can be rated. In the first statement a score of 1 means that you think almost all of the people in that group are "intelligent." A score of 7 means that you think almost all of the people in the group are "unintelligent." A score of 4 means that you think the group is not towards one end or the other, and of course you may choose any number in between that comes closest to where you think people in the group stand.

1. Where would you rate whites in general on this scale?

2. Where would you rate blacks in general on this scale?

3. Where would you rate Asian Americans in general on this scale?

4. Where would you rate Hispanic Americans in general on this scale?

Im going to show you a second a seven-point scale on which the characteristics of people in a group can be rated. In the first statement a score of 1 means that you think almost all of the people in that group are trustworthy. A score of 7 means that you think almost all of the people in the group are untrustworthy. A score of 4 means that you think the group is not towards one end or the other, and of course you may choose any number in between that comes closest to where you think people in the group stand.

1. Where would you rate whites in general on this scale?

2. Where would you rate blacks in general on this scale?

3. Where would you rate Asian Americans in general on this scale?

4. Where would you rate Hispanic Americans in general on this scale?

---

[1]Please note the item order for each question is randomized.

Interest in Politics

1. In general, how interested are you in politics?

2. During a typical week, how many days do you watch the national network evening news programs on television (e.g., ABC, CBS, CNN, FOX, NBC)?

3. During a typical week, how many days do you read about the national and local news in a daily newspaper (not including sports and entertainment)?

4. During a typical week, how many days do you watch, read, or listen to news on the Internet (not including sports and entertainment)?

5. How often do you discuss politics with your family and friends?

Political knowledge items

The items used to construct the political knowledge scale are as follow:

1. How long is the term of office for someone elected to the U.S House of Representatives?

2. Do you happen to know which party controls (holds a majority of seats in) the House of Representatives at the moment?

3. Who has the final responsibility of deciding whether a law is constitutional or not?

4. Do you recall the president of Russia?

5. Do you recall who is the Secretary of Defense?

6. Do you recall who is the Secretary of State?

7. Do you know which party controls (holds a majority of seats in) the U.S. Senate at the moment?

8. How long is the term of office for someone elected to the U.S. Senate? class.

Policy Issue Stances

1. In the 2012 U.S. presidential election, did you vote for Barack Obama, Mitt Romney, some other candidate, or were you unable to get to the polls that day?

2. We hear a lot of talk these days about liberals and conservatives. Here is a scale on which the political views that people might hold are arranged from extremely liberal to extremely conservative. Where would you place yourself upon this scale, or haven't you thought much about this?

3. Some people think the government should provide many more services, even if it means an increase in spending. Suppose these people are at one end of the scale, at point 1. Other people believe that government should provide fewer services, even in areas such as health and education, in order to reduce spending. Suppose these people are at the other end of the scale, at point 7. And of course other people fall somewhere in between these two opposite views. Where would you place yourself on this scale?

4. Some people feel that the government in Washington should see to it that every person has a job and a good standard of living. Suppose these people are at one end of a scale, at point 1. Others think the government should just let each person get ahead on their own. Suppose these people are at the other end, at point 7. Where would you put yourself on this scale?

5. In 2010 the government passed comprehensive health care reform that is designed to provide many more Americans with access to affordable health care, and to help contain the rapidly rising cost of health insurance. How strongly do you support or oppose this health care reform?

6. Some people believe that we should spend much less money for defense. Others feel that defense spending should be greatly increased. Where would you place yourself on this scale?

7. Some people believe the United States should solve international problems by using diplomacy and other forms of international pressure, and use military force only if absolutely necessary. Suppose we put such people at point 1 on this scale. Others believe that diplomacy and pressure often fail, and the U.S. must be ready to use military force. Suppose we put them at point 7. And of course other people fall somewhere in between these two endpoints. Where would you put yourself on this scale?

8. How important of a goal should combating international terrorism be for the country?

9. In 2012, President Obama signed an agreement with Afghan President Hamid Karzai to handover combat operations to Afghan forces by spring 2013. It is expected that full U.S. withdrawal from Afghanistan will take place by the end of 2014. Some people believe the additional troops have only served to get more Americans killed, and we should begin withdrawing all of our troops from that country immediately. Others believe that it is so important to strengthen the government of Afghanistan, oppose the Taliban, and fight Al Qaeda that even

more troops should be sent to the country. And, of course, others want to stick with President Obama's original plan. What is your opinion on this issue?

10. There has been much discussion in recent years about abortion. The opinions below represent many perspectives on this issue. Which one best agrees with your view?

11. Early in his presidency, George Bush passed the No Child Left Behind Act, which employed standardized tests to hold school districts accountable for the federal aid they were receiving. Do you favor or oppose this set of educational reforms?

12. Do you favor or oppose a school voucher program that would allow parents to use tax funds to send their children to the school of their choice, even if it were a private school?

13. Some people say that because of past discrimination, African Americans should be given preference in hiring and promotion. Others say that such preference in hiring and promotion of blacks is wrong because it gives blacks advantages they haven't earned. How strongly do you agree with either of these opinions?

14. The United States is currently highly dependent upon importing oil from the Middle East and other countries around the world to meet our energy needs. Lately, people have begun talking about energy independence. One way to accomplish this is to begin using alternative energy sources, such as nuclear, solar, and hydroelectric sources of energy. Another option is to produce more oil domestically. Which of these options do you believe is more important?

15. Do you favor or oppose leasing federally owned lands, such as in national parks and wilderness areas, to oil companies for exploration and drilling?

16. Should federal spending on dealing with crime be increased, decreased or kept about the same?

17. Should federal spending on aid to poor people – that is, welfare – be increased, decreased or kept about the same?

18. How strongly do you agree or disagree that there should be a two-year limit on how long someone can receive welfare benefits?

19. In recent years, warnings from the scientific community have increased regarding global warming. According to these sources, the accumulation of greenhouse gases in the earth's atmosphere is leading to changes in the planet's weather patterns. Do you favor or oppose the government taking steps to combat global warming?

20. It is estimated that in recent years, the number of illegal immigrants living in the United States has surpassed 12 million. Some policymakers are calling for stricter measures to reduce these numbers, including the building of a fence along the border with Mexico, and forced deportation of anyone found to be living here illegally. Do you favor or oppose stricter laws concerning illegal immigration?

21. Do you favor or oppose the use of capital punishment – the death penalty – in the criminal justice system?

22. Should the federal government do more, do less, or keep its current policies toward protecting the environment?

23. Taking everything into account, do you think the war in Iraq has been worth the cost or not?

24. In recent years, several states have passed or attempted to pass laws legalizing same-sex marriage. Do you favor or oppose giving lesbians and homosexuals the right to marry?

25. Do you favor or oppose stricter gun control laws in your state?

26. In 2001, the Bush administration passed a series of tax cuts that lowered income tax rates for many Americans, although the greatest savings went to those at the very top of the income scale. In December 2010 Congress extended those tax cuts for another two years, even for the richest Americans. Do you favor or oppose the continuation of these tax cuts?

8.2 Appendix for Chapter 4

Sample descriptives

The following set of tables provides the descriptives across a number of my observables.

Gender

Item	Count	%
Male	111	48
Female	122	52
Total	233	100

Race

Item	Count	%
White	141	61
Black	16	7
East Asian	15	6
South Asian	25	11
Hispanic	22	9
Other	14	6
Total	233	100

Party Identification

Item	Count	%
Moderate/DK	38	16
Democrat	131	56
Republican	64	27
Total	233	100

Major

Item	Count	%
History	9	4
Mathematics	1	0
Economics	15	6
Political Science	121	52
Engineering	8	3
Biology	14	6
Psychology	5	2
Other	60	26
Total	233	100

Family Income

Item	Count	%
Less than 10k	8	3
10 to 20k	8	3
20 to 30k	14	6
30 to 40k	16	7
40 to 50k	18	8
50 to 75k	32	14
75 to 100k	42	18
100 to 150k	53	23
150k or more	42	18
Total	233	100

Figure A4.1: Density plot: % of total opened African American candidate across explicit bias
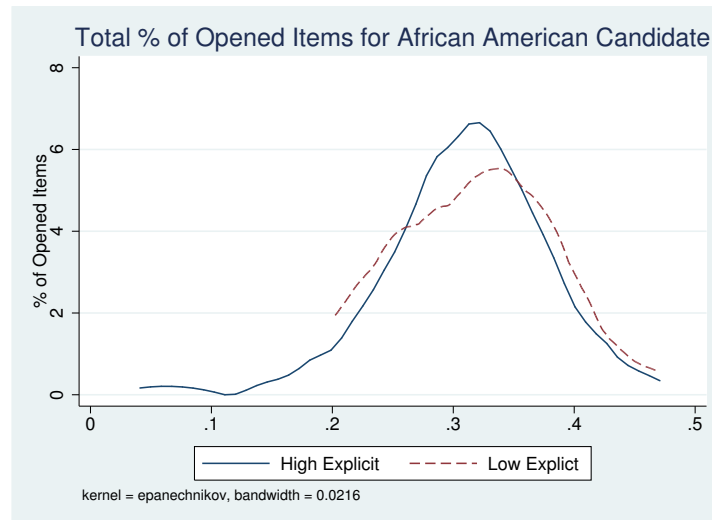
Total % of Opened Items for African American Candidate

kernel = epanechnikov, bandwidth = 0.0216

A Kolmogorov-Smirnov test for equality of distribution functions is rejected at the .05 level

Figure A4.2: Density plot: % of reading time African American candidate across explicit bias
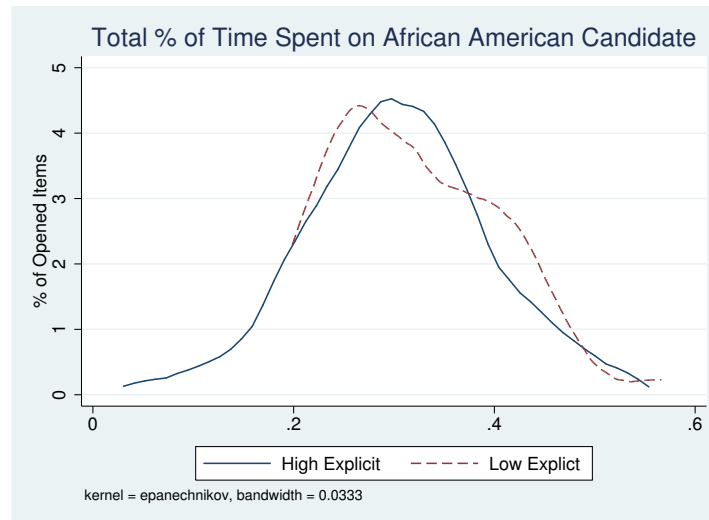
A Kolmogorov-Smirnov test for equality of distribution functions is rejected at the .05 level

Regression models with alternative IAT specification

Here I report all of the results of the regression models using an alternative method to compute IAT scores. Specifically, I compute IAT scores here by computing the median block score and then differencing B3, B4 from B6, B7. For simplicity, I only report the coefficients of main interest. The models, which include the coefficients for the covariates are available upon request.

Table A4.1: Regression: % of opened information for African American candidate

	(1)	(2)	(3)
IAT	-0.0001**	-0.0001**	-0.0001*
	(0.000)	(0.000)	(0.000)
Explicit Bias		0.0004	0.0001
		(0.002)	(0.002)
Total Opened			0.023
			(0.018)
Controls	No	No	Yes
Constant	0.328***	0.327***	0.248***
	(0.007)	(0.008)	(0.089)
N	115	115	115
R^2	0.046	0.047	0.110

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]The dependent variable is calculated as a percentage. Specifically it is calculated by dividing the total amount of information opened for the African American candidate by the total amount of information opened for all candidates.

[b]The IAT variable is constructed by computing the median block score and then differencing B3, B4 from B6, B7.

[c]The Conservative and Liberal dummies were constructed from the standard ideology scale given to subjects. The baseline category is moderate.

Table A4.2: Regression: % of reading time for African American candidate

	(1)	(2)	(3)	(4)
IAT	-0.0001**	-0.0001**	-0.0001**	-0.0001**
	(0.000)	(0.000)	(0.000)	(0.000)
Explicit Bias		0.001	0.001	0.001
		(0.003)	(0.003)	(0.003)
Reading Speed			-0.012	-0.009
			(0.018)	(0.020)
Total Duration				0.020
				(0.029)
Controls	No	No	Yes	Yes
Constant	0.329***	0.327***	0.425***	0.283
	(0.011)	(0.011)	(0.125)	(0.256)
N	115	115	115	115
R^2	0.041	0.043	0.087	0.090

Standard errors in parentheses

[*] $p < 0.10$, [**] $p < 0.05$, [***] $p < 0.01$

[a]The reading speed variable was calculated by dividing the subjects' time spent on reading the three sets of directions prior to the experiment by the total number of words for the directions.

Table A4.3: Regression: Candidate Evaluation, Including Info Processing Variables

	(1)	(2)	(3)
IAT	-0.014	-0.018†	-0.006
	(0.015)	(0.014)	(0.016)
Explicit Bias	-1.204	-1.051	-0.849
	(1.119)	(1.200)	(0.964)
Memory Count African American		3.734***	1.066
		(1.297)	(1.539)
Memory Count White Moderate		-4.392***	-1.776†
		(1.023)	(1.224)
% Opened African American			1.507***
			(0.410)
% Opened White Moderate			-0.691**
			(0.351)
Controls	Yes	Yes	Yes
Constant	13.945	12.482	-18.957
	(12.232)	(11.891)	(20.815)
N	115	114	114
R^2	0.08	0.20	0.311

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]The memory count variables are constructed from the memory recall tasks that subjects were given at the end of the experiment. Each unique item listed for a given candidate was coded as one memory. The accuracy of memories have not yet been coded.

[b]The % Opened African American variable is the percentage of opened items that were about the African American candidate relative to the total amount opened.

Table A4.4: Logistic Regression: Vote for African American candIdate

	(1)	(2)	(3)
IAT	-0.002	-0.003*	-0.001
	(0.001)	(0.001)	(0.002)
Explicit Bias	-0. 105	-0.160	-0.318
	(0.153)	(0.188)	(0.252)
Memory Count African American		0.777***	1.062***
		(0.195)	(0.337)
Memory Count White Moderate		-0.558***	-0.890***
		(0.169)	(0.321)
% Opened African American			0.278***
			(0.078)
% Opened White Moderate			0.149*
			(0.088)
Controls	Yes	Yes	Yes
Constant	-0.970	-2.628✝	-17.291***
	(1.296)	(1.679)	(4.750)
N	115	114	114
R^2	0.083	0.285	0.501

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]The outcome variable here is 1 if subject voted for the African American and 0 if they voted for the other candidates or abstained.

Regression models with alternative specification for explicit bias

In this section, I report the results from my models where I replace the explicit bias measure used in the paper with the Kinder and Sears (1982, 2004) symbolic racism scale.

Table A4.5: Regression: % of opened information for African American candidate

	(1)	(2)	(3)
IAT	-0.075**	-0.07**	-0.065**
	(0.032)	(0.031)	(0.033)
Explicit Bias		-0.005	-0.006
		(0.007)	(0.006)
Total Opened			0.033*
			(0.018)
Controls	No	No	Yes
Constant	0.330***	0.345***	0.225***
	(0.009)	(0.022)	(0.083)
N	115	115	115
R^2	0.037	0.043	0.117

Standard errors in parentheses

$^*\ p < 0.10$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$

[a]The dependent variable is calculated as a percentage. Specifically it is calculated by dividing the total amount of information opened for the African American candidate by the total amount of information opened for all candidates.

[b]The IAT variable is constructed by computing the log transformed block score and then differencing B3, B4 from B6, B7.

[c]The Conservative and Liberal dummies were constructed from the standard ideology scale given to subjects. The baseline category is moderate.

Table A4.6: Regression: % of reading time for African American candidate

	(1)	(2)	(3)	(4)
IAT	-0.093**	-0.083*	-0.079*	-0.075✝
	(0.047)	(0.045)	(0.047)	(0.047)
Explicit Bias		-0.012	-0.010	-0.012
		(0.008)	(0.009)	(0.009)
Reading Speed			-0.009	-0.003
			(0.018)	(0.020)
Total Duration				0.042
				(0.031)
Controls	No	No	Yes	Yes
Constant	0.331***	0.363***	0.443***	0.155
	(0.013)	(0.029)	(0.126)	(0.256)
N	115	115	115	115
R^2	0.031	0.048	0.088	0.099

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]The reading speed variable was calculated by dividing the subjects' time spent on reading the three sets of directions prior to the experiment by the total number of words for the directions.

Table A4.7: Regression: Candidate Evaluation, Including Info Processing Variables

	(1)	(2)	(3)
IAT	-23.719	-27.111✝	-18.468
	(18.076)	(17.281)	(17.996)
Explicit Bias	-0.670	-0.476	1.468
	(3.042)	(2.861)	(2.841)
Memory Count African American		3.969***	1.208
		(1.293)	(1.527)
Memory Count White Moderate		-4.405***	-1.745✝
		(1.032)	(1.185)
% Opened African American			1.458***
			(0.405)
% Opened White Moderate			-0.794**
			(0.359)
Controls	Yes	Yes	Yes
Constant	12.333	7.153	-21.282
	(14.684)	(13.811)	(22.449)
N	115	114	114
R^2	0.081	0.204	0.315

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]The memory count variables are constructed from the memory recall tasks that subjects were given at the end of the experiment. Each unique item listed for a given candidate was coded as one memory. The accuracy of memories have not yet been coded.

[b]The % Opened African American variable is the percentage of opened items that were about the African American candidate relative to the total amount opened.

Table A4.8: Logistic Regression: Vote for African American candidate

	(1)	(2)	(3)
IAT	-2.342†	-3.590*	-3.113
	(1.658)	(1.930)	(2.606)
Explicit Bias	-0.019	0.110	0.099
	(0.278)	(0.322)	(0.374)
Memory Count African American		0.786***	0.975***
		(0.195)	(0.306)
Memory Count White Moderate		-0.552***	-0.788***
		(0.170)	(0.290)
% Opened African American			0.281***
			(0.079)
% Opened White Moderate			0.123†
			(0.084)
Controls	Yes	Yes	Yes
Constant	-1.228	-3.445*	-17.571***
	(1.457)	(1.932)	(4.892)
N	115	114	114
R^2	0.081	0.288	0.495

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]The outcome variable here is 1 if subject voted for the African American and 0 otherwise.

8.3   Appendix for Chapter 5

Sample descriptives

The following set of tables provides the crosstabs for the experimental conditions across a number of my observables. For all of our demographic variables, I report no significant imbalances across conditions.

Gender by condition

Condition	Male	Female
Control	59	56
Treatment	52	66
Total	111	122

Party identification by condition

Condition	Democrat	Republican	Independent
Control	67	31	17
Treatment	64	33	21
Total	131	64	38

Race by condition

Condition	White	Black	Hispanic	South Asian	East Asian	Other
Control	71	5	12	11	7	9
Treatment	70	10	11	14	8	5
Total	141	16	22	25	15	14

Regression models with alternative IAT specification

Here I report all of the results of the regression models using an alternative method to compute IAT scores. Specifically, I compute IAT scores here by computing the median block score and then differencing B3, B4 from B6, B7. For simplicity, I only report the coefficients of main interest. The models, which include the coefficients for the covariates are available upon request.

Table A5.1: Regression: % of Information Opened for African American Candidate (Across Conditions)

	(1)	(2)	(3)
IAT	-0.000*	-0.0001***	-0.0001***
	(0.000)	(0.000)	(0.000)
Explicit Bias	-0.003**	0.0004	0.0001
	(0.001)	(0.002)	(0.002)
Treatment		0.001	-0.006
		(0.012)	(0.012)
Treatment * IAT		0.0001**	0.0001**
		(0.000)	(0.000)
Treatment * Explicit		-0.006**	-0.007***
		(0.002)	(0.003)
Controls	No	No	Yes
Constant	0.329***	0.327***	0.353***
	0.006	(0.008)	(0.077)
N	233	233	233
R^2	0.033	0.062	0.116

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]Controls include gender, race, political knowledge, partisanship, and the total number of opened items during the primary

Table A5.2: Regression: % of Reading Time for African American Candidate (Across Conditions)

	(1)	(2)	(3)
IAT	-0.0001*	-0.0001***	-0.0001***
	(0.032)	(0.000)	(0.000)
Explicit Bias	-0.004**	0.001	0.001
	(0.002)	(0.003)	(0.003)
Treatment		0.001	-0.0003
		(0.016)	(0.016)
Treatment * IAT		0.0001*	0.0001*
		(0.065)	(0.069)
Treatment * Explicit		-0.009***	-0.010***
		(0.003)	(0.004)
Controls	No	No	Yes
Constant	0.329***	0.327***	0.293
	0.008	(0.011)	(0.218)
N	233	233	233
R^2	0.032	0.063	0.110

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]Controls include gender, race, political knowledge, partisanship, and the total number of opened items during the primary

Table A5.3: Regression: Candidate Evaluation, Including Info Processing Variables

	(1)	(2)	(3)
IAT	-0.013†	-0.023*	-0.021†
	(0.008)	(0.014)	(0.014)
Explicit Bias	-1.819***	-1.215	-1.064
	(0.650)	(1.093)	(1.176)
Treatment		-3.179	-6.600†
		(4.842)	(4.657)
Treatment * IAT		0.020	0.026†
		(0.018)	(0.017)
Treatment * Explicit		-1.071	-0.829
		(1.316)	(1.315)
Controls	No	No	Yes
Constant	4.366**	6.042†	9.515
	2.249	(4.166)	(8.354)
N	233	233	231
R^2	0.051	0.057	0.275

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]Controls include gender, race, political knowledge, partisanship, and memory count

Table A5.4: Logistic Regression: Vote for African American Candidate

	(1)	(2)	(3)
IAT	-0.001**	-0.002*	-0.003*
	(0.881)	(0.001)	(0.001)
Explicit Bias	-0.152**	-0.104	-0.223
	(0.071)	(0.131)	(0.173)
Treatment		0.213	-0.045
		(0.439)	(0.516)
Treatment * IAT		0.001	0.001
		(0.002)	(0.002)
Treatment * Explicit		-0.075	0.093
		(0.157)	(0.201)
Controls	No	No	Yes
Constant	-0.708***	-0.830***	-0.379
	0.208	(0.347)	(1.082)
N	233	233	231
R^2	0.035	0.054	0.273

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]Controls include gender, race, political knowledge, partisanship, and the total number of opened items during the primary
[b]The outcome variable here is 1 if subject voted for the African American and 0 if they voted for the other candidates, or abstained.

Regression models with alternative specification for explicit bias

In this section, I report the results from my models where I replace the explicit bias measure used in the paper with the Kinder and Sears (1982, 2004) symbolic racism scale.

Table A5.5: Regression: % of Information Opened for African American Candidate (Across Conditions)

	(1)	(2)	(3)
IAT	-0.006†	-0.071**	-0.006*
	(0.004)	(0.031)	(0.032)
Explicit Bias	-0.027	-0.005	-0.006
	(0.023)	(0.007)	(0.006)
Treatment		0.002	0.003
		(0.030)	(0.030)
Treatment * IAT		0.080*	0.086*
		(0.045)	(0.046)
Treatment * Explicit		-0.003	-0.004
		(0.009)	(0.009)
Controls	No	No	Yes
Constant	0.343***	0.345***	0.336***
	0.015	(0.022)	(0.076)
N	233	233	233
R^2	0.016	0.031	0.083

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]Controls include gender, race, political knowledge, partisanship, and the total number of opened items during the primary

Table A5.6: Regression: % of Reading Time for African American Candidate (Across Conditions)

	(1)	(2)	(3)
IAT	-0.027	-0.083*	-0.084*
	(0.032)	(0.045)	(0.047)
Explicit Bias	-0.013***	-0.012	-0.012
	(0.006)	(0.008)	(0.008)
Treatment		-0.001	-0.004
		(0.038)	(0.038)
Treatment * IAT		0.100†	0.111*
		(0.062)	(0.065)
Treatment * Explicit		-0.003	-0.003
		(0.012)	(0.012)
Controls	No	No	Yes
Constant	0.359***	0.363***	0.190
	0.019	(0.029)	(0.223)
N	233	233	233
R^2	0.027	0.041	0.083

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]Controls include gender, race, political knowledge, partisanship, and the total number of opened items during the primary

Table A5.7: Regression: Candidate Evaluation, Including Info Processing Variables

	(1)	(2)	(3)
IAT	-12.551	-30.430*	27.834*
	(9.663)	(16.862)	(16.693)
Explicit Bias	-3.118*	-1.191	0.168
	(1.876)	(2.748)	(2.769)
Treatment		3.717	-3.262
		(10.309)	(8.864)
Treatment * IAT		31.089†	32.187*
		(20.446)	(19.228)
Treatment * Explicit		-3.906	-2.512
		(3.709)	(3.315)
Controls	No	No	Yes
Constant	12.341**	11.051	8.030
	4.937	(7.954)	(10.374)
N	233	233	231
R^2	0.027	0.041	0.260

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]Controls include gender, race, political knowledge, partisanship, and memory count

Table A5.8: Logistic Regression: Vote for African American Candidate

	(1)	(2)	(3)
IAT	-1.642[*]	-2.936[**]	-3.169[*]
	(0.878)	(1.524)	(1.755)
Explicit Bias	-0.138	0.031	-0.075
	(0.157)	(0.245)	(0.293)
Treatment		0.938	-0.152
		(0.980)	(1.176)
Treatment * IAT		2.217	2.160
		(1.868)	(2.144)
Treatment * Explicit		-0.320	0.0002
		(0.321)	(0.380)
Controls	No	No	Yes
Constant	-0.390	-0.859	-0.619
	0.474	(0.760)	(1.347)
N	233	233	231
R^2	0.019	0.036	0.262

Standard errors in parentheses
[*] $p < 0.10$, [**] $p < 0.05$, [***] $p < 0.01$

[a]Controls include gender, race, political knowledge, partisanship, and the total number of opened items during the primary
[b]The outcome variable here is 1 if subject voted for the African American and 0 if they voted for the other candidates, or abstained.

8.4 Appendix for Chapter 6

Table A6.1: The impact of affective state on total duration

	All	African American	White Moderate	White Partisan
IAT	-30.939	-27.564✝	1.534	16.659
	(46.385)	(19.891.)	(22.247)	(28.025)
Explicit Bias	11.170	-5.942✝	3.946	1.695
	(9.760)	(3.996)	(4.994)	(5.579)
Affective State	-0.385	-0.280	0.643*	-0.385
	(0.804)	(0.445)	(0.385)	(0.610)
Controls	Yes	Yes	Yes	Yes
Constant	403.003***	-559.046***	-579.663	-561.558
	(46.112)	(93.538)	(104.900)	(107.781)
N	115	115	115	115
R^2	0.034	0.403	0.331	0.240

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]Controls for this model specification include gender, race, political knowledge, partisanship, the total number of opened items, the total reading time, and reading speed

Table A6.2: The impact of emotional state on total number of opened items (Using Ethnocentrism Scale)

	All	African American	White Moderate	White Partisan
IAT	7.370	-4.086✝	-1.871	8.079**
	(15.272)	(3.045)	(3.354)	(4.446)
Explicit Bias	-0.9108	-0.308	0.523***	-0.162
	(1.008)	(0.300)	(0.177)	(0.283)
Emotional state	-0.132	-0.079✝	0.026	-0.016
	(0.303)	(0.055)	(0.045)	(0.074)
Controls	Yes	Yes	Yes	Yes
Constant	91.567***	-103.866	-97.056	-96.937
	(12.532)	(9.796)	(10.025)	(11.733)
N	115	115	115	115
R^2	0.027	0.711	0.6831	0.587

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]All models were run using robust standard errors. Controls for models include gender, race, political knowledge, partisanship, the total number of opened items (only included in the candidate specific models), the total reading time, and reading speed

Table A6.3: Information processing for African American candidate, across emotional state (Using Ethnocentrism Scale)

	Total Opened	Total Duration
IAT	-0.065*	-0.082✝
	(0.035)	(0.050)
Explicit Bias	-0.002	-0.001
	(0.003)	(0.005)
Emotional state	-0.000	-0.0005
	(0.001)	(0.001)
Emotional state * IAT	-0.001	-0.003
	(0.003)	(0.004)
Emotional state*Explicit	0.001✝	0.0001
	(0.0005)	(0.0003)
Controls	Yes	Yes
Constant	0.197**	0.106
	(0.090)	(0.196)
N	115	115
R^2	0.125	0.088

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

[a]All models were run using robust standard errors. Controls for models include gender, race, political knowledge, partisanship, the total number of opened items, the total reading time, and reading speed

Family income by condition

Condition	Under 40k	41-50k	50-75k	75-100k	100-150k	More than 150K
Control	29	10	16	14	26	20
Treatment	17	8	16	28	27	22
Total	46	18	32	42	53	42

Major by condition

Condition	Political Science	History	Economics	Engineering	Biology	Other
Control	56	4	6	5	9	35
Treatment	65	5	9	3	5	31
Total	121	9	15	8	14	66

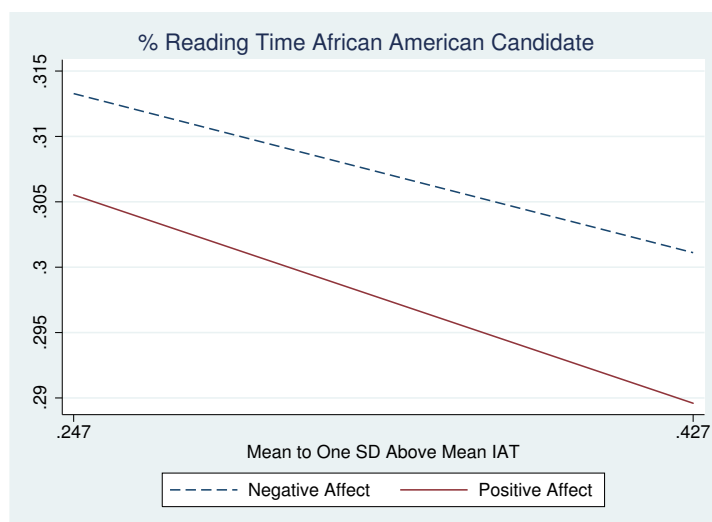Figure A6.1: IAT effect across affective state scale (total duration)

Figure A6.2: Symbolic racism effect across affective state scale (total duration)
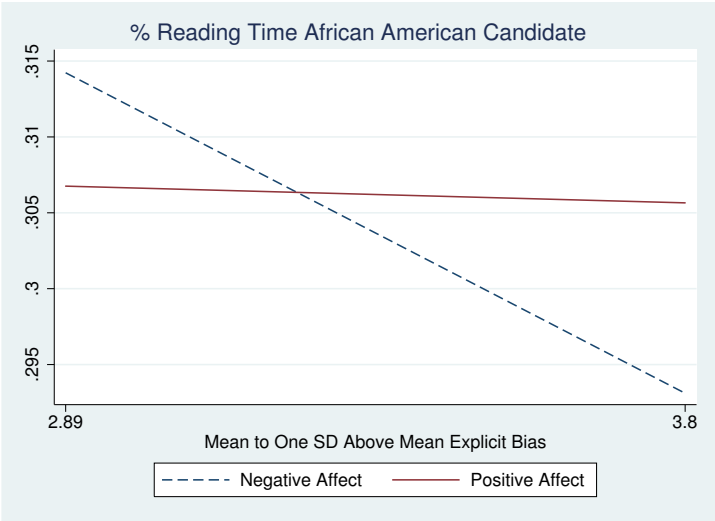

Figure A6.3: Marginal effect of IAT across affective state scale(total duration)
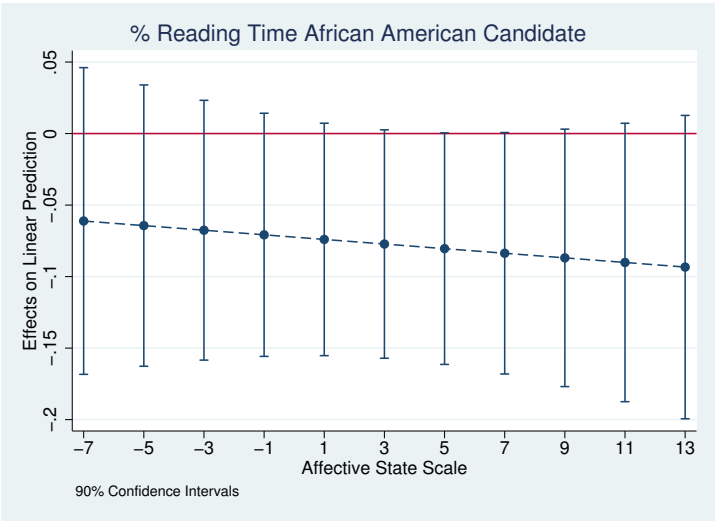
Figure A6.4: Marginal effect of symbolic racism across affective state scale(total duration)
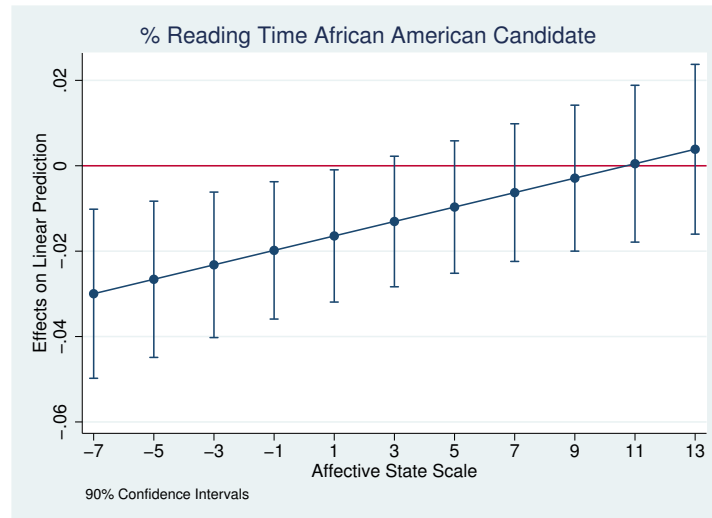

Figure A6.5: Symbolic racism effect at different levels of affective state, Total duration African American candidate
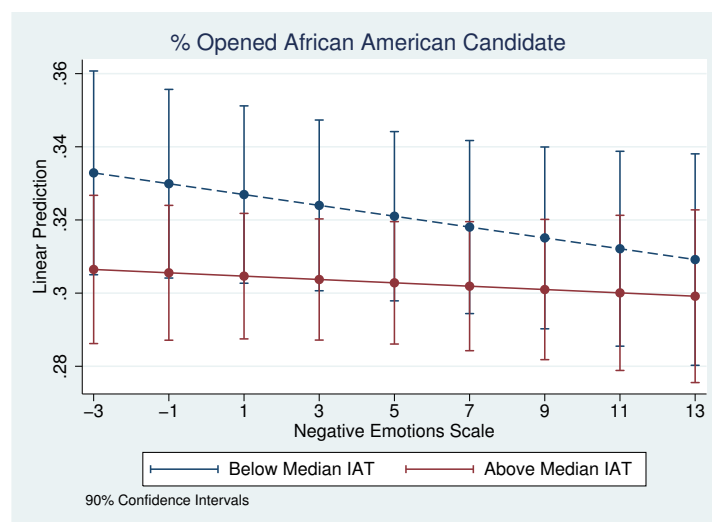
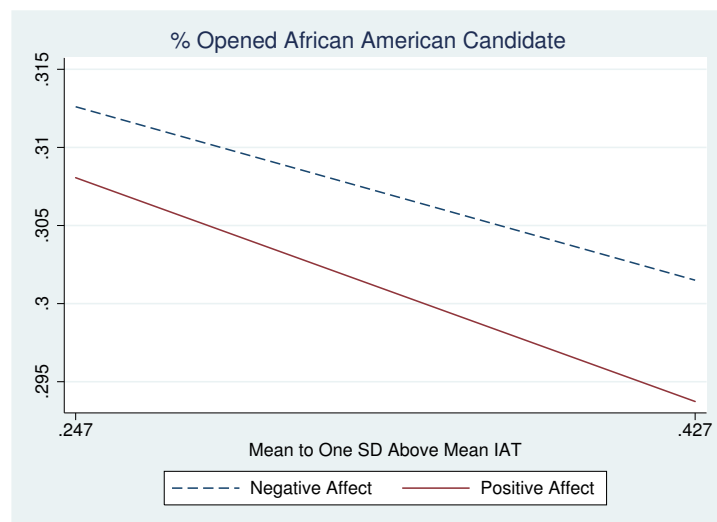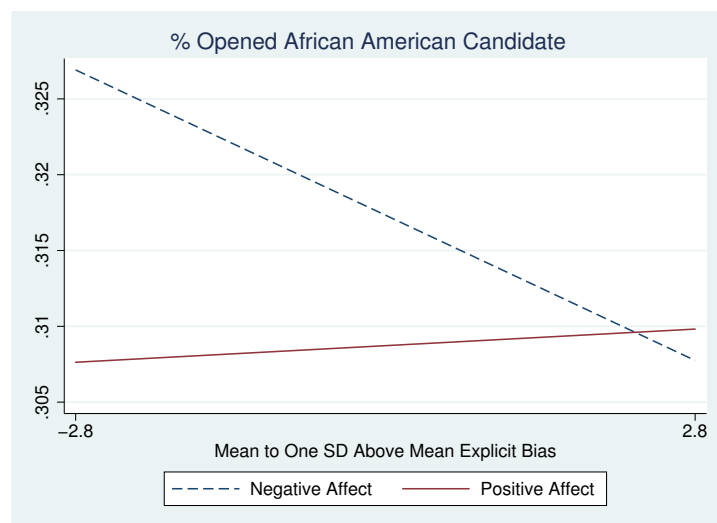Figure A6.6: IAT effect across emotional state scale (Using Ethnocentrism Scale)

Figure A6.7: Ethnocentrism scale effect across emotional state scale

Bibliography

[1] Albarracin, D., & Kumkale, G. T. 2003. "Affect as information in persuasion: A model of affect identification and discounting." *Journal of Personality and Social Psychology*, 29, 834-845.

[2] Abelson, R.P., & Levi, A.1985."Decision Making and Decision Theory" in *The Handbook of Social Psychology Vol 1. 3rd Ed.*, ed. Gardner, L. & Aronson, E. New York: Random House

[3] Albertson, B. 2008. "Religious Appeals and Implicit Attitudes" Paper presented at the annual meeting of the Midwest Political Science Association, Chicago IL

[4] Anderson, J. 1983. *The Architecture of Cognition.* Cambridge, MA: Harvard University Press

[5] Anderson, J., Bothell, D., Bryne, M., Douglass, S., Lebiere, C., & Qin, Y. 2004. "An Integrated Theory of Mind." *Psychological Review* 111(4): 1036-1060.

[6] Ansolabehere, S., & Iyengar., S. 1997. *Going Negative* Simon & Schuster press

[7] Arcuri, L., Castelli L., Silvia G., Zogmaister C., & Amadori. A. 2008. "Predicting the Vote: Implicit Attitudes as Predictors of the Future Behavior of Decided and Undecided Voters." *Political Psychology* 29(3): 369-387.

[8] Arkes, H. R., & Tetlock, P. E. 2004. "Attributions of implicit prejudice, or" would Jesse Jackson 'fail' the implicit association test?". *Psychological Inquiry*, 15(4), 257-278.

[9] Asendorpf, J. B., Banse, R., & Schnabel, K. 2006. "Employing automatic approach and avoidance tendencies for the assessment of implicit personality self-concept." *Experimental Psychology*, 53(1), 69-76.

[10] Asendorpf, J. B., Banse, R., & Mucke, D. 2002. "Double dissociation between implicit and explicit personality self-concept: The case of shy behavior." *Journal of Personality and Social Psychology*, 83, 380-393.

[11] Atkinson, R.C & Shiffrin, R.M. 1968. "Chapter: Human memory: A proposed system and its control processes" in Spence, K.W. & Spence, J.T. *The psychology of learning and motivation (Volume 2).* New York: Academic Press. pp. 89-195.

[12] Ballew, C. C., & Todorov, A. 2007. "Predicting political elections from rapid and unreflective face judgments." *Proceedings of the National Academy of Sciences*,104, 17948-17953.

[13] Banaji et al. 2012. *Project Implicit* website

[14] Banaji, M. R.,& Hardin, C. D. 1996. "Automatic stereotyping." *Psychological Science*, 7, 136-141.

[15] Banaji, M.R. & Greenwald, A.R. 1994. "Implicit Stereotyping and Prejudice" in *The Psychology of Prejudice: The Ontario Symposium*, eds. Zanna, M.P & Olson, J.M., 55-76, Hillsdale NJ: Erlbaum.

[16] Banse, R., Seise, J., & Zerbes, N. 2001. "Implicit attitudes toward homosexuality: Reliability, validity, and controllability of the IAT.". *Zeitschrift fur Experimentelle Psychologie*, 48, 145-160.

[17] Bar-Anan, Y., & Nosek, B. A. 2012. "A comparative investigation of seven implicit measures of social cognition." Unpublished manuscript.

[18] Bargh., J. et al. 1992. "The generality of the automatic attitude activation effect." *Journal of Personality and Social Psychology* Vol. 62(6) 893-912

[19] Bargh, J. A. 1994. "The four horseman of automaticity: Awareness, intention, efficiency and control in social cognition." In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition: Vol. 1. Basic processes* (2nd ed., pp. 1-40). Hillsdale, NJ: Erlbaum.

[20] Bargh, J.A. 1997. "The Automaticity of Everyday Life" in *The Automaticity of Everyday Life: Advances in Social Cognition* edited by R.S. Wyer. Mahwah: Eribaum

[21] Bargh, J. 1999. *The Cognitive Monster: The Case Against the Controllability of Automatic Stereotype Effects. In Dual-Process Theories in Social Psychology*, ed. Shelley Chaiken Yaacov Trope. The Guilford Press.

[22] Bartels, L. M. 2000. "Partisanship and voting behavior, 1952-1996." *American Journal of Political Science*, 44, 35-50.

[23] Batson, C. Daniel, Turk, Cynthia L., Shaw, Laura L., Klein, & Tricia R. *Journal of Personality and Social Psychology*, Vol 68(2), Feb 1995, 300-313.

[24] Bertrand, M. et al. 2005. "Implicit Discrimination." *American Economic Review* Vol. 94(4), 991-1013

[25] Berinsky, A. et al. 2012. "Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk." *Political Analysis* Vol. 20, 351-368.

[26] Betsch, T., Plessner, H., Schwieren, C., & Gutig, R. 2001. "I like it but I dont know why: A value-account approach to implicit attitude formation." *Personality and Social Psychology Bulletin*, 27, 242- 253.

[27] Blair, I.V. & Banaji, M.R. 1996. "Automatic and Controlled Processes in Stereotype Priming." *Journal of Personality and Social Psychology* 70: 1142-63.

[28] Blanton, H., Jaccard, J., Gonzales, P. M., & Christie, C. 2006. "Decoding the implicit association test: Implications for criterion prediction." *Journal of Experimental Social Psychology*, 42(2), 192-212.

[29] Blanton, H., Jaccard, J. 2006. "Arbitrary metrics in psychology." *American Psychologist*, 61(1), 27.

[30] Blanton, H. & J. Jaccard. 2008. "Unconscious Racism: A Concept in Pursuit of a Measure" *Annual Review of Sociology* Vol. 34, 277-97

[31] Bodenhausen, G. V., Kramer, G. P., & Ssser, K. 1994. "Happiness and stereotypic thinking in social judgment." *Journal of Personality and Social Psychology*, 66(4), 621.

[32] Bodenhausen, G. V., Sheppard, L. A., & Kramer, G. P. 1994. "Negative affect and social judgment: The differential impact of anger and sadness." *European Journal of Social Psychology*, 24(1), 45-62.

[33] Bosson, J. K., Swann Jr, W. B., & Pennebaker, J. W. 2000. "Stalking the perfect measure of implicit self-esteem: The blind men and the elephant revisited?" *Journal of Personality and Social Psychology*, 79(4), 631.

[34] Brader, T. 2005. "Striking a responsive chord: How political ads motivate and persuade voters by appealing to emotions." *American Journal of Political Science*, 49(2), 388-405.

[35] Brader, T. 2006. *Campaigning for Hearts and Minds: How Emotional Appeals in Political Ads Work*. University of Chicago Press.

[36] Brader, T. 2008. "What Triggers Public Opposition to Immigration? Anxiety, Group Cues, and Immigration Threat." *American Journal of Political Science* Vol. 52(4) 959-978.

[37] Bravo, J. & Greco. V. unpublished manuscript. "You Get What You Pay For? Incentivizing Subjects as a Means to Improve the Measurement Properties of the IAT"

[38] Burdein, Lodge, M., & Taber, C. 2004. "Implicit identifications in political information processing: An experimental test of the hot identification hypothesis." Unpublished manuscript, Stony Brook University.

[39] Burdein, I., Lodge. M & Taber. C. 2006. "Experiments on the Automaticity of Political Beliefs and Attitudes." *Political Psychology* 27(3):359-371.

[40] Claxton, G. 1997. *Hare brain, tortoise mind: How intelligence increases when you think less.* New York: Harper Collins.

[41] Carraro, L. & Castelli.L. 2010. "The Implicit and Explicit Effects of Negative Political Campaigns: Is the Source Really Blamed?" *Political Psychology*, 31 (4): 617-645.

[42] Carraro, L., Gawronski, B., & Castelli, L. 2010. "Losing on All Fronts: The Effects of Negative versus Positive Person-Based Campaigns on Implicit and Explicit Evaluations of Political Candidates." *British Journal of Social Psychology*, 49 (3): 453-470.

[43] Clore, G. & Isbell, L. 2001. "Emotion and Virtue and Vice" in *Political Psychology in Practice,* ed. J. Kuklinski, 103-123, New York: Cambridge University Press.

[44] Collins, A. & Loftus, E. 1975. "A Spreading-Activation Theory of Semantic Processing." *Psychological Review* 82 (6): 407-428.

[45] Cosmides, L., & Tooby, J. 2000, "Evolutionary Psychology and the Emotions," in *Handbook of Emotions*, ed. Michael Lewis and Jeannette M. Haviland-Jones, New York: Guilford, 91-115.

[46] Cunningham, W. A., Preacher, K. J., & Banaji, M. R. 2001. "Implicit attitude measures: Consistency, stability, and convergent validity." *Psychological Science*, 12(2), 163-170.

[47] Cunningham, W., Raye, L., & Johnson, K. 2004. "Implicit and explicit evaluation: fMRI correlates of valence, emotional intensity, and control in the processing of attitudes." *Journal of Cognitive Neuroscience*, 16, 1717-1729.

[48] Damasio, A.R. 1994. *Descartes' Error: Emotion, Reason, and the Human Brain.* New York: Putnam

[49] Damasio, A.R. 1996. "The Somatic Marker Hypothesis and the Possible Functions of the Prefrontal Cortex." *Proceedings of the Royal Society* 351: 1413-1420.

[50] Dasgupta, N., Desteno, D., Williams, L., Hunsinger, M. 2009. "Fanning the Flames of Prejudice: The Influence of Specific Incidental Emotions on Implicit Prejudice." *Emotion*, 9 (4): 585-591.

[51] Dasgupta, N., & Rivera, L. M. 2006. "From automatic anti-gay prejudice to behavior: The moderating role of conscious beliefs about gender and behavioral control." *Journal of Personality and Social Psychology*, 91, 268-280.

[52] De Houwer J.& Moors A. 2005. "How to define and examine the implicitness of implicit measures." In *Implicit Measures of Attitudes*, ed. B Wittenbrink, N Schwarz, pp. 179-94. New York: Guilford

[53] DeSteno, D., Dasgupta, N., Bartlett, M. Y., & Cajdric, A. 2004. "Prejudice from thin air: The effect of emotion on automatic intergroup attitudes." *Psychological Science*, 15, 319-324.

[54] Dijksterhuis, A. 2004b. "Think different: The merits of unconscious thought in preference development and decision making." *Journal of Personality and Social Psychology*, 87, 586-598.

[55] Dijksterhuis,A., Bos, M.W., Nordgren,L.F., & van Baaren, R.B.2006. "On making the right choice: The deliberation-without-attention effect." *Science*, 311, 1005-1007.

[56] Dijksterhuis, A., & Meurs, T. 2006. "Where creativity resides: The generative power of unconscious thought." *Consciousness and Cognition*, 15, 135-146.

[57] Dijksterhuis, A., & van Olden, Z. 2006. "On the benefits of thinking unconsciously: Unconscious thought can increase post-choice satisfaction." *Journal of Experimental Social Psychology*, 42, 627-631.

[58] Dijksterhuis, A. 2009. "Automaticity and the Unconscious" in *Handbook of Psychology* edited by Susan Fiske, Daniel Gilbert, and Gardner Lindzey, John Wiley and Sons, 227-267.

[59] Ditonto, T., Lau, R., & Sears, D. 2013. "AMPing Racial Attitudes: Comparing the Power of Explicit and Implicit Racism Measures in 2008." *Political Psychology* 34 (4) 487-510.

[60] Dovidio, J. F., Kawakami, K., Johnson, C., Johnson, B., & Howard, A. 1997. "On the nature of prejudice: Automatic and controlled processes." *Journal of Experimental Social Psychology*, 33, 510-540.

[61] Dovidio, J. F., Kawakami, K., & Gaertner, S. L. 2002. "Implicit and explicit prejudice and interracial interaction." *Journal of Personality and Social Psychology*, 84, 754-770.

[62] Egloff, B., & Schmukle, S. C. 2002. "Predictive validity of an Implicit Association Test for measuring anxiety." *Journal of Personality and Social Psychology*, 83, 1441-1455.

[63] Egloff, B., Schwerdtfeger, A., & Schmukle, S. C. 2005. "Temporal stability of the Implicit Association Test anxiety." *Journal of Personality Assessment*, 84(1), 82-88

[64] Epstein, S. "Cognitive-Experiential Self-Theory of Personality," in *Handbook of Psychology, Vol. 5, Personality and Social Psychology*, ed. Theodore Millon, Melvin J. Lerner, and Irving B. Weiner, Hoboken, NJ: Wiley, 159-84.

[65] Erisen, C., Lodge, M., & Taber, C. 2013. "Affective Contagion in Effortful Political Thinking" *Political Psychology*

[66] Evans, J. 2003. "In Two Minds: Dual-Process Accounts of Reasoning." *TRENDS in Cognitive Sciences* 7(10):454-459.

[67] Evans, J. 2008. "Dual-Processing Accounts of Reasoning, Judgement, and Social Cognition." *Annual Review of Psychology* 59: 255-278.

[68] Evans, J. 2010. *Thinking Twice: Two Minds in One Brain.* Oxford University Press.

[69] Evans, J., & Stanovich, K. 2013. "Dual-Process Theories of Higher Cognition: Advancing the Debate." *Perspectives on Psychological Science* 8 (3): 223-241.

[70] Fazio, R.H. 2007. "Attitudes as Object-Evaluation Association of Varying Strength." *Social Cognition*, 25 (5): 603-637.

[71] Fazio, R. H., & Olson, M. A. 2003. "Implicit measures in social cognition research: Their meaning and use." *Annual Review of Psychology*, 54(1), 297-327.

[72] Fazio, R.H., & Towles-Schwen, T. 1999. "The MODE model of attitude- behavior processes" in eds. S. Chaiken & Y. Trope, *Dual process theories in social psychology,* 97116. New York: Guilford Press.

[73] Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. 1995. "Variability in automatic activation as an unobtrusive measure of racial attitudes. A bona fide pipeline?" *Journal of Personality and Social Psychology*, 69, 1013-1027.

[74] Fazio, R. H. 1990. "Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework." In M. P. Zanna (Ed.), *Advances in experimental social psychology* Vol. 23, pp. 75?109. San Diego: Academic Press.

[75] Fazio, R.H. et al., 1986. "On the Automatic Activation of Attitudes." *Journal of Experimental Social Psychology* Vol. 18, 339-57

[76] Fazio. R.H., Chen, J., McDonel, E.C., & Sherman, S.J. 1982. "Attitude Accessibility, attitude-behavior consistency, and the strength of the object-evaluation association." *Journal of Experimental Social Psychology*, 18, 339-357.

[77] Feldman B, L., & Russell, J. A. 1999. "The structure of current affect: Controversies and emerging consensus". *Psychological Science*, 8, 10-14.

[78] Ferguson, M. & Porter, S. 2010. "What is Implicit about Goal Pursuit?" in *Handbook of Implicit Social Cognition: Measurement, Theory, and Applications*, eds. B. Gawronski & K. Payne, 311-331. New York: Guilford.

[79] Finn, C. & Glaser, J. 2010. "Voter Affect and the 2008 U.S. Presidential Election: Hope and Race Mattered" *Analyses of Social Issues and Public Policy*, 10: 262-275.

[80] Finn, C. & Glaser, J. 2013. "How and Why Implicit Attitudes Should Affect Voting" *PS: Political Science and Politics*, Vol 46, No. 3, 537-44

[81] Friese, M., Smith, C. T., Plischke, T., Bluemke, M., & Nosek, B. A. 2012. "Do implicit attitudes predict actual voting behavior particularly for undecided voters?" *PLoS ONE*, 7(8), e44130.

[82] Friese, M., Hofmann, W., & Wnke, M. 2009. "The impulsive consumer: Predicting consumer behavior with implicit reaction time measures." *Social Psychology and Consumer Behavior*, 335-64.

[83] Friese, M., Hofmann, W., & Schmitt, M. 2008. "When and why do implicit measures predict behaviour? Empirical evidence for the moderating role of opportunity, motivation, and process reliance." *European Review of Social Psychology*, 19, 285-338.

[84] Friese, M., Bluemke, M., Wnke, M. 2007. "Predicting voting behavior with implicit attitude measures: The 2002 German parliamentary election." *Experimental Psychology*, 54, 247-255.

[85] Frijda, N. H. 2009. "Emotion experience and its varieties". *Emotion Review*, 1, 264-271.

[86] Galdi, S., Gawronski, B., Arcuri, L., & Friese, M. 2012. "Selective Exposure in Decided and Undecided Individuals: Differential Relations to Automatic Associations and Conscious Beliefs." *Personality and Social Psychology Bulletin*, 38 (5): 559-569.

[87] Gawronksi, B., Galdi, S., & Arcuri, L. 2013. "What Can Political Psychology Learn from Implicit Measures? Empirical Evidence and New Directions." *Political Psychology*, in press.

[88] Gibson, B. 2008. "Can evaluative conditioning change attitudes toward mature brands? New evidence from the Implicit Association Test." *Journal of Consumer Research*, 35(1), 178-188.

[89] Gilens, M. 1999. *Why Americans Hate Welfare.* Chicago: University of Chicago Press.

[90] Glaser, J., & Knowles, E. D. 2008. "Implicit motivation to control prejudice." *Journal of Experimental Social Psychology*, 44, 164-172.

[91] Green, A. R., Carney, D. R., Pallin, D. J., Ngo, L. H., Raymond, & K. L., Iezzoni, L. 2007. "Implicit bias among physicians and its prediction of thrombolysis decisions for black and white patients." *Journal of General Internal Medicine*, 22, 1231-1238.

[92] Greenwald, A. G., & Nosek, B. A. 2001. "Health of the Implicit Association Test at age 3." *Zeitschrift fu?r Experimentelle Psychologie*, 48, 85-93.

[93] Greenwald, A.G. & Banaji, M.R. 1995. "Implicit Social Cognition: Attitudes, Self-Esteem, and Stereotypes." *Psychological Review* 4-27.

[94] Greenwald, A.G., Draine, S.C., & Abrams, R.L. 1996. "Three Cognitive Markers of Unconscious Semantic Activation." *Science* 273: 1699-1702.

[95] Greenwald, A.G et al. 1998. "Measuring individual differences in implicit cognition: The implicit association test." *Journal of Personality and Social Psychology* Vol. 74, 1464-1480

[96] Hastie, R., & Dawes R. 2001. *Rational Choice in an Uncertain World*, Sage Publications

[97] Hassin, R., Uleman, J., & Bargh, J. 2005. *The New Unconscious*. New York, Oxford University Press.

[98] Hedrick, J. & Ksiazkiewicz, A. 2013. "Implicit Attitudes in Political Science Research." *PS: Political Science and Politics* Vol 46, No. 3, 525-551.

[99] Henry, P. J., & Sears, D. O. 2002. "The Symbolic Racism 2000 scale". *Political Psychology*, 23, 253-283.

[100] Hofmann, W., Gawronski, B., Gschwendner, T., Le, H., & Schmitt, M. 2005. "A meta-analysis on the correlation between the Implicit Association Test and explicit self-report measures." *Personality and Social Psychology Bulletin*, 31(10), 1369-1385

[101] Huddy, L., & Feldman, S. . 2009. "On Assessing the Political Effects of Racial Prejudice." *Annual Review of Political Science*, 12, 423-47.

[102] Isbell, L. M., Ottati, V. C., & Burns, K. C. 2006. "Affect and politics: Effects on judgment, processing, and information seeking." In D. P. Redlawsk (Ed.), *Feeling politics* (pp. 5786). New York, NY: Palmgrave.

[103] Jellison, W.A. et al. 2004. "Implicit and explicit measures of sexual orientation attitudes: In group preferences and related behaviors and beliefs among gay and straight men." *Personality and Social Psychology Bulletin* Vol. 30, 629-642

[104] Jost, J., Rudman, L., Blair, I., Carney, D., Dasgupta, N., Glaser, J., & Hardin, C. 2009. "The existence of implicit bias is beyond reasonable doubt: A refutation of ideological and methodological objections and executive summary of ten studies that no manager should ignore." *Research in Organizational Behavior* 29, 39-69.

[105] Jost, J. T., Banaji, M. R., & Nosek, B. A. 2004. "A decade of system justification theory: Accumulated evidence of conscious and unconscious bolstering of the status quo." *Political Psychology*, 25, 881-919

[106] Kahneman, D. 2011. *Thinking, fast and slow.* New York, NY: Farrar, Straus and Giroux.

[107] Kahneman D, Slovic P, & Tversky A. 1982. *Judgment Under Uncertainty: Heuristics and Biases.* Cambridge, UK: Cambridge Univ. Press

[108] Kahneman D, & Tversky A. 1972. "Subjective probability: a judgment of representativeness." *Cognitive Psychology.* 3: 430-54

[109] Kahneman D, & Tversky A. 1973. "On the psychology of prediction." *Psychological Review* 80: 237-51

[110] Kam, C. D., & Kinder, D. R. 2007. "Terror and ethnocentrism: Foundations of American support for the war on terrorism." *Journal of Politics*, 69, 320-338.
Kam, C.D. 2007. "Implicit Attitudes, Explicit Choices: When Subliminal Priming Predicts Candidate Preference." *Political Behavior* 29(3): 343-367.

[111] Kim, D. Y. 2003. "Voluntary controllability of the Implicit Association Test (IAT)." *Social Psychology Quarterly*, 66, 83-96.

[112] Kinder, D. & Ryan, T. unpublished manuscript. "Prejudice and Politics Re-Examined: The Political Significance of Implicit Racial Bias."

[113] Kinder, D., & Dale-Riddle, A. 2012. *The End of Race? Obama, 2008, and Racial Politics in America.* New Haven, CT: Yale University Press.

[114] Kinder, D., and Sanders, L. 1996. *Divided By Color.* Chicago: University of Chicago Press.

[115] Kinder, D. R., & Sears, D. O. 1981. "Prejudice and politics: Symbolic racism versus racial threats to the good life." *Journal of Personality and Social Psychology*, 40, 414-431.

[116] Kleinberg, M., Ditonto, T., Andersen. D, & Lau. R "Speeders, Cheaters, and Satisficers. Mechanical Turk: Attention in Online vs. Laboratory Experiments." In progress.

[117] Knoll, B. "Implicit Nativist Attitudes, Social Desirability, and Immigration Policy Preferences." *International Migration Review*, 47 (1): 132-165.

[118] Krosnick, J. A., Lupia, A. & DeBell, M. 2010. "Methodology Report and Users Guide for the 20082009 ANES Panel Study." Palo Alto, CA, and Ann Arbor, MI: Stanford University and the University of Michigan.

[119] Kuklinski, J. H., & Quirk, P. J. 2000. "Reconsidering the rational public: Cognition, heuristics, and mass opinion." In A. Lupia, M. D. McCubbins, S. L. Popkin (Eds.), *Elements of reason: Cognition, choice, and the bounds of rationality* (pp. 153182). Cambridge: Cambridge University Press.

[120] Lau, R., & Redlawsk, D. 1997. "Voting correctly." *American Political Science Review*, 91(3), 585-598.

[121] Lau R. et al.1999. " The Effects of Negative Political Advertisements: A Meta-Analytic Assessment" *American Political Science Review*, Vol. 93(4)

[122] Lau R. et al. 2007. " The Effects of Negative Political Advertisements: A Meta-Analytic Reassessment" *The Journal of Politics*, Vol. 69(4)

[123] Lau, Richard R. & Redlawsk .D. 2006. *How Voters Decide: Information Processing During Election Campaigns*. New York, NY: Cambridge University Press.

[124] LeDoux, J. 1994. "Emotion, Memory, and the Brain." *Scientific American*, 270 (6): 32-39.

[125] LeDoux, J. 1996. *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*. New York: Simon & Schuster

[126] Lee, L., Amir, On., & Ariely, D. 2009. "In Search of Homo Economicus: Cognitive Noise and the Role of Emotion in Preference Consistency." *Journal of Consumer Research*, Vol. 36, No. 2, 173-187.

[127] Lerner, J. S., & Keltner, D. 2001. "Fear, anger, and risk." *Journal of Personality and Social Psychology*, 81, 146-159.

[128] Lerner, J. S., & Tiedens, L. Z. 2006. "Portrait of the angry decision maker: How appraisal tendencies shape anger's influence on cognition." *Journal of Behavioral Decision Making*, 19(2), 115-137.

[129] Levine, G.M., Halberstadt, J.B., & Goldstone, R.L. 1996. "Reasoning and the weighting of attributes in attitude judgments." *Journal of Personality and Social Psychology*, 70, 230-240.

[130] Lodge, M., Taber, C. S. 2000. "Three steps toward a theory of motivated political reasoning." In A. Lupia, M. McCubbins, S. Popkin (Eds.), *Elements of Reason: Cognition, Choice, and the Bounds of Rationality*, Cambridge: Cambridge University Press.

[131] Lodge, M. & Taber. C. 2005. "The Automaticity of Affect for Political Leaders, Groups, and Issues: An Experimental Test Employing." *Political Psychology* 26:455-482.

[132] Lodge, M. & Taber. C. 2013. *The Rationalizing Voter*, Cambridge University Press

[133] Lyle, M.L. 2008. "Dueling Consciousness: The Psychological Implications of the American Ideology of Racial Dominance" *Doctoral Dissertation* Department of Political Science, Duke University

[134] Mackuen, M., Wolak, J., Keele, L., & G. Marcus. 2010. "Civic Engagements: Resolute Partisanship or Reflective Deliberation." *American Journal of Political Science*, vol 54(2), 440-458.

[135] Maholtra. N et al. 2013 "Economic Explanations for Opposition to Immigration: Distinguishing Between Prevalence and Magnitude." *American Journal of Political Science* forthcoming

[136] Marcus, G. & MacKuen, M. 1993. "Anxiety, enthusiasm, and the vote: The emotional underpinnings of learning and involvement during presidential campaigns." *American Political Science Review*, 87(September), 672-685.

[137] Marcus, G., Newman, R., & MacKuen, M. 2000. '*Affective intelligence and political judgment.* Chicago: University of Chicago Press.

[138] Marcus, G. E., & MacKuen, M. B. 2001. "Emotions and politics: The dynamic functions of emotionality." *Citizens and politics: Perspectives from political psychology*, 41-67.

[139] Marcus, G. 2003. "The Psychology of Emotion and Politics." In *Oxford Handbook of Political Psychology*, edited by David Sears, Leonie Huddy, and Robert Jervis. New York: Oxford University Press, 182-221.

[140] McConnell, A. R, & Leibold, J. M. 2001. "Relations among the Implicit Association Test, explicit attitudes, and discriminatory behavior." *Journal of Experimental Social Psychology*, 37, 435-442.

[141] McConahay, J.B. 1986. "Modern Racism, Ambivalence, and the Modern Racism Scale" in *Prejudice, Discrimination, and Racism*, eds. Dovidio, J.F. & Gaertner, S.L, 91-126, New York: Academic Press

[142] Mendleberg, T. 2001. *The Race Card*, Princeton NJ: Princeton University Press

[143] Merolla, J. & Zechmeister. E. 2009. "Terrorist Threat, Leadership, and the Vote: Evidence from Three Experiments." *Political Behavior*, 31: 575-601.

[144] Miller, G.A. 1956. "Information Theory." *Scientific American*, 195, 42-46.

[145] Miller, G.A. 1957. "The Magic Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information." *Psychological Review* 63 (2): 81-97.

[146] Mo. C. 2011. "Dual Process Theories of the Mind and Vote Choice: The Consequences of Implicit and Explicit Attitudes on the Judgment of Voters." *Doctoral Dissertation,* Graduate School of Business, Stanford University

[147] Morris, J., Squires, N., Taber, C., Lodge, M. 2003. "Activation of Political Attitudes: A Psychophysiological Examination of the Hot Cognition Hypothesis." *Political Psychology*, 24 (4): 727-745.

[148] Neely, J.H. 1977. "Semantic Priming and Retrieval from Lexical Memory: Roles of Inhitbitionless Spreading Activation and Limited Capacity Attention." *Journal of Experimental Psychology* 106: 226-254.

[149] Neisser, U. 1976. *Principles and Implications of Cognitive Psychology.* New York: W.H. Freeman.

[150] Niederle, M. & Gneezy, U. 2007. "Do Women Shy Away from Competition? Do Men Compete Too Much?" *Quarterly Journal of Economics*, Vol. 122, No. 3, 1067-1101.

[151] Nordgren, L.F., & Dijksterhuis, A. 2006. "Conscious thought and decisional noise." *Unpublished manuscript*, University of Amsterdam, Amsterdam, The Netherlands.

[152] Norretranders, T. 1998. *The User Illusion: Cutting Consciousness Down to Size.* New York: Penguin Books

[153] Nosek, B. A., Greenwald, A. G., & Banaji, M. R. 2005. "Understanding and using the Implicit Association Test: II. Method variables and construct validity." *Personality and Social Psychology Bulletin*, 31(2), 166-180.

[154] Nosek, B. A., Greenwald, A. G., & Banaji, M. R. 2007. "The Implicit Association Test at age 7: A methodological and conceptual review." In J. A. Bargh (Ed.), *Social psychology and the unconscious: The automaticity of higher mental processes* (pp. 265-292). New York, NY: Psychology Press.

[155] Nosek, B. A., Banaji, M. R., & Greenwald, A. G. 2002. "Harvesting implicit group attitudes and beliefs from a demonstration web site." *Group Dynamics: Theory, Research, and Practice*, 6, 101-115.

[156] Nosek, B. 2009. "Race Attitude Measures Predicted in the 2008 U.S. Presidential Election." *Analyses of Social Issues and Public Policy* Vol. 9, 241-253

[157] Olivola, C. Y., & Todorov, A. 2010. "Elected in 100 milliseconds: Appearance-based trait inferences and voting." *Journal of Nonverbal Behavior.*

[158] Paas, F., Renkl, A., & Sweller, J. 2004. "Cognitive load theory: Instructional implications of the interaction between information structures and cognitive architecture." *Instructional science*, 32(1), 1-8.

[159] Parker, M.T., & Isbell, L. 2010. "How I Vote Depends on How I Feel: The Differential Impact of Anger and Fear on Political Information Processing." *Psychological Science* 21 (4): 548-550.

[160] Pasek, J., et al. 2009. "Determinants of Turnout and Candidate Choice in the 2008 U.S. Presidential Election." *Public Opinion Quarterly* 73: 943-994.

[161] Payne, J., Samper, A., Bettman, J., Luces, M. 2008. "Boundary Conditions on Unconscious Thought in Complex Decision Making." *Psychological Science*, Vol. 19, 11.

[162] Payne, B. Keith. 2001. "Prejudice and Perception: The Role of Automatic and Controlled Processes in Misperceiving a Weapon." *Journal of Personality and Social Psychology* 81:181-192.

[163] Payne, B.K., Cheng, S. M., Govorun, O., & Stewart B. D. 2005. "An inkblot for attitudes: Affect misattribution as implicit measurement. *Journal of Personality and Social Psychology*, 89, 277-293.

[164] Payne, K. et al. 2010. "Implicit and Explicit Prejudice in the 2008 Presidential Election." *Journal of Experimental Social Psychology* 46: 367-374.

[165] Peffley, M., & Hurwitz, J. 2007. "Persuasion and resistance: Race and the death penalty in America." *American Journal of Political Science* 51 (4): 996-1012.

[166] Peffley, M., Hurwitz, J., & Sniderman, P. 1997. "Racial stereotypes and whites political views of blacks in the context of welfare and crime." *American Journal of Political Science* 41: 30-60.

[167] Perez, E. 2010. "Explicit Evidence on the Implicit Attitude of Immigration." *Political Behavior* 32(4): 517-545.

[168] Perez. E. 2013. "Implicit Attitudes: Meaning, Measurement, and Synergy with Political Science." *Politics, Groups, and Identities* 1(2): 275-297.

[169] Perugini, M., & Leone, L. 2004. "Individual differences in moral decision making: Validation of an implicit measure of morality." Unpublished manuscript.

[170] Perugini, M., Richetin, J., & Zogmaister, C. 2010. "Implicit Social Cognition and Indirect Measures in Consumer Behavior." in *Handbook of Implicit Social Cognition, Measurement, Theory, and Applications*, eds. B. Gawronski & K. Payne, 255-277, New York: Guilford.

[171] Pham, M., Cohen, J., Pracejus., J & Hughes, D. 2001, "Affect Monitoring and the Primacy of Feelings in Judgment," *Journal of Consumer Research*, 28, 167- 88.

[172] Pham, Michel T. 2004, "The Logic of Feeling," *Journal of Consumer Psychology*, 14 (4), 360-69.

[173] Phelps, E. A., O'Connor, K. J., Cunningham, W. A., Gatenby, J. C., Funayama, E. S., & Gore, J. C. 2000. "Amygdala activation predicts performance on indirect measures of racial bias." *Journal of Cognitive Neuroscience*, 12, 729-738.

[174] Quattrone, G. A., & Tversky, A. 1988. "Contrasting rational and psychological analysis of political choice." *American Political Science Review*, 82, 716-736.

[175] Ratcliff, R. & McKoon, G. 1988. "A Retrieval Theory of Priming in Memory." *Psychological Review*, Vol. 95 (3): 385-408.

[176] Ratcliff, R. & McKoon, G. 1994. "Retrieving information from memory: Spreading-activation theories versus compound-cue theories." *Psychological Review*, Vol. 101 (1): 177-184.

[177] Ratcliff, R. & McKoon, G. 1996. "Bias Effects in Implicit Memory Tasks." *Journal of Experimental Psychology* 125 (4): 403-421.

[178] Redlawsk, D. 2001. "You must remember this: A test of the on-line voting model." *Journal of Politics*, 63, 29-58.

[179] Redlawsk, D. 2002. "Hot cognition or cool consideration? Testing the effects of motivated reasoning on political decision making." *Journal of Politics*, 64, 1021-1044.

[180] Redlawsk, D. 2010. "The Affective Tipping Point: Do Motivated Reasoners ever Get It?" *Political Psychology* 31(4): 563-593

[181] Rooth, D. 2010. "Automatic Associations and Discrimination in Hiring: Real World Evidence." *Labour Economics* Vol. 17, 523-534

[182] Riggio, H., & Riggio, R. 2010. "Appearance-Based Trait Inferences and Voting: Evolutionary Roots and Implications for Leadership." *Journal of Nonverbal Behavior*, Vol. 34, 119-125.

[183] Rosch, E. 1975. "Cognitive Representations of Semantic Categories." *Journal of Experimental Psychology* 104: 192-233.

[184] Rudman, L. 2011. *Implicit Measures for Social and Personality Psychology.* Newbury Park, CA: Sage.

[185] Rudman, L. A., & Ashmore, R. D. 2007. "Discrimination and the Implicit Association Test." *Group Processes and Intergroup Relations*, 10, 359-372.

[186] Rudman, L. A., & Glick, P. 2001. "Prescriptive gender stereotypes and backlash toward agentic women." *Journal of Social Issues*, 57, 743-762.

[187] Rudman, L. A., & Lee, M. R. 2002. "Implicit and explicit consequences of exposure to violent and misogynous rap music." *Group Processes and Intergroup Relations*, 5, 133-150.

[188] Sanbonmatsu, D. M., & Fazio, R. H. 1990. "The role of attitudes in memory?based decision making." *Journal of Personality and Social Psychology*, 59, 614?622.

[189] Shiffrin, R.M., & Schneider, W. 1977. "Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory." *Psychological Review* Vol 84(2), 127-190

[190] Schmukle, S. C., & Egloff, B. 2004. "Does the Implicit Association Test for assessing anxiety measure trait and state variance?" *European Journal of Personality*, 18, 483-494.

[191] Schmukle, S. C., & Egloff, B. 2005. "A latent state-trait analysis of implicit and explicit personality measures". *European Journal of Psychological Assessment*, 21(2), 100-107.

[192] Schooler, J.W., Ohlsson, S., & Brooks, K. 1993. "Thoughts beyond words: When language overshadows insight." *Journal of Experimental Psychology: General*, 122, 166-183.

[193] Schooler, J.W., & Melcher, J. 1995. "The ineffability of insight." In S.M. Smith, T.B. Ward, R.A. Finke (Eds.), *The creative cognition approach* (pp. 97134). Cambridge, MA: MIT Press.

[194] Schuette, R. A., & Fazio, R. H. 1995. "Attitude accessibility and motivation as determinants of biased processing: A test of the MODE model." *Personality and Social Psychology Bulletin*, 21, 704?710.

[195] Sears, D.O.1988."Symbolic Racism." in *Eliminating Racism: Profiles in Controversy* ,eds. P. Katz, & Dalmas A. Taylor. New York: Plenum.

[196] Segura, G. & Valenzuela, A. 2010. "Hope, Tropes, and Dopes: Hispanic and White Racial Animus in the 2008 Election." *Presidential Studies Quarterly* 40 (3) 497-514.

[197] ShIffrin, R. M., & Schneider, W. 1977. "Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory." *Psychological Review*, 84, 127-190.

[198] Shiv., B. & Fedorikhin, A. 1999. " Heart and Mind in Conflict: The Interplay of Affect and Cognition in Consumer Decision Making." *Journal of Consumer Research* Vol 26. 278-292.

[199] Sidanius, J., Pratto, F., & Bobo, L. 1996. "Racism, Conservatism, Affirmative Action, and Intellectual Sophistication: A Matter of Principled Conservatism or Group Dominance?" *Journal of Personality and Social Psychology*, 70 (March): 476-490.

[200] Simon, H.A. 1957. *Models of Man.* New York, NY: Wiley.

[201] Simon, H.A. 1979. "Rational decision making in business organizations." *American Economic Review*, 69, 493-513.

[202] Slovic, P., Finucane, M., Peters, E., & McGregor, D. 2004. "Risk as Analysis and Risk as Feelings: Some Thoughts about Affect, Reason, Risk, and Feelings." *Risk Analysis*, 24 : 1 - 12.

[203] Slovic, P., Finucane, M., Peters, E. & McGregor, D. 2007. "The Affect Heuristic." *European Journal of Operations Research*, 177 (3): 1333 - 1352.

[204] Slovic, P. 1999. "Trust, Emotion, Sex, Politics, and Science: Surveying the Risk Assessment Battlefield." *Risk Analysis*, 19: 689-701.

[205] Smith. C., & Nosek., B. 2011. " Affective focus increases the concordance between implicit and explicit attitudes." *Social Psychology* Vol. 42(4), 300-313.

[206] Sniderman, P., Brody, R., Tetlock, P. 1991. *Reasoning and Choice: Explorations in Political Psychology*, Cambridge University Press.

[207] Spence., A. & Townsend. E. 2008. "Spontaneous evaluations: Similarities and differences between affect heuristic and implicit attitudes." *Cognition & Emotion* Vol. 22(1)

[208] Steffens, M. C. 2004. "Is the Implicit Association Test immune to faking?" *Experimental Psychology*, 51(3), 165-179.

[209] Stanovich KE. 1999. *Who is Rational? Studies of Individual Differences in Reasoning.* Mahwah, NJ: Elrbaum

[210] Stanovich, K. E. 2011. *Rationality and the reflective mind.* New York, NY: Oxford University Press.

[211] Stanovich, K. & West. R. 2000. "Individual Differences in Reasoning: Implications for the Rationality Debate." *Behavioral and Brain Sciences* 23(5):645-665.

[212] Steffens, M. C., & Buchner, A. 2003. "Implicit Association Test: Separating transsituationally stable and variable components of attitudes toward gay men." *Experimental Psychology*, 50(1), 33-48.

[213] Strack, F., & Deutsch, R. 2004. "Reflective and impulsive determinants of social behavior." *Personality and Social Psychology Review*, 8(3), 220-247.

[214] Swann, W. 1990. "To be adored or to be known? The interplay of self-enhancement and self-verification" in eds. Higgins, E. & Sorrentino, R. *Handbook of Motivation and Cognition: Foundations of Social Behavior*, Vol. 2. , (pp. 408-448). New York, NY, US: Guilford Press,

[215] Taber, C. S., & Lodge, M. 2006. "Motivated skepticism in the evaluation of political beliefs." *American Journal of Political Science*, 50(3), 755-769

[216] Taylor, S. 1981. "A Categorization Approach to Stereotyping" in *Cognitive Processes in Stereotyping and Intergroup Behavior,* ed. D. Hamilton, 145-182, Hillsdale, NJ: Erlbaum.

[217] Tesler, M. & Sears, D.O. 2010. *Obamas race. The 2008 election and the dream of a post-racial America.* Chicago: University of Chicago Press.

[218] Tetlock, P. E., & Mitchell, G. 2009. "Implicit Bias and Accountability Systems: What Must Organizations do to Prevent Discrimination?" *Research in Organizational Behavior*, 29, 3-38.

[219] Todorov, A., Said, C. P., Engell, A. D., Oosterhof, N. N. 2008. "Understanding evaluation of faces on social dimensions." *Trends in Cognitive Sciences*, 12, 455460.

[220] Todorov, A., Mandisodza, A. N., Goren, A., & Hall, C. 2005. "Inferences of competence from faces predict election outcomes." *Science*, 308, 1623-1626.

[221] Tsetsos, K. Chater, N., & Usher, M. 2012. "Salience driven value integration explains decision biases and preference reversal." *Proceedings of the National Academy of Sciences*, 109, no. 24.

[222] Tversky A, & Kahneman D. 1973. "Availability: a heuristic for judging frequency and probability." *Cognitive Psychology.* 5: 207-32

[223] Tversky, A., & Kahneman, D. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science*, Vol. 185, No. 4157, 1124-1131.

[224] Valentino, N., Hutchings, V., Banks, A., & Davis, A. 2008. "Is a Worried Citizen a Good Citizen? Emotions, Political Information Seeking, and Learning via the Internet." *Political Psychology* 29 (2): 247-273

[225] Verhulst, B., Lodge, M., & H. Lavine. 2010. "The Attractiveness Halo: Why Some Candidates are Perceived More Favorably than Others." *Journal of Nonverbal Behavior*, v34, p111-117.

[226] von Hippel, W, Brener, L., & von Hippel, C. 2008. "Implicit prejudice toward injecting drug users predicts intentions to change jobs among drug and alcohol nurses." *Psychological Science*, 19, 7-11.

[227] Watson, D., & Clark, L. A. 1999. The PANAS-X: Manual for the positive and negative affect schedule-expanded form.

[228] Westling., E., Mann, T., & Ward., A. 2006. "Self-Control of Smoking: When Does Narrowed Attention Help?" *Journal of Applied Social Psychology* 36(9) 2115-2133.

[229] Wilson, T.D., & Schooler, J.W. 1991. "Thinking too much: Introspection can reduce the quality of preferences and decisions." *Journal of Personality and Social Psychology*, 60, 181-192.

[230] Wilson, T.D., Lisle, D., Schooler, J.W., Hodges, S.D., Klaaren, K.J., & LaFleur, S.J. 1993. "Introspecting about reasons can reduce post- choice satisfaction." *Personality and Social Psychology Bulletin*, 19, 331-339.

[231] Wilson, T.D. 2002. '*Strangers to ourselves: Discovering the adaptive unconscious.* Cambridge, MA: Harvard University Press.

[232] Winter, N., & Berinsky. A. 1999. "What's Your Temperature? Thermometer Ratings and Political Analysis." Presented at the Annual Meeting of the American Political Science Association, Atlanta, GA.

[233] Zajonc, R. 1980. "Feeling and Thinking: Preferences Need No Inferences." *American Psychologist* 39(2) 117-123.

[234] Zajonc. R. 2000. "Feeling and Thinking: Closing the Debate Over the Independence of Affect" in *Feeling and Thinking: The Role of Affect in Social Cognition*, ed. J. Forgas, 31-58, Cambridge, MA: Cambridge University Press

[235] Ziegert, J. C., & Hanges, P. J. 2005. "Employment discrimination: The role of implicit attitudes, motivation, and a climate for racial bias." *Journal of Applied Psychology*, 90, 553-562.