

IMPLICIT AND EXPLICIT DETERMINANTS OF JUDGMENT AND BEHAVIOUR:

A descriptive and prescriptive analysis

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

In Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Jeremy David Cone

August 2012

© 2012 Jeremy David Cone

IMPLICIT AND EXPLICIT DETERMINANTS OF JUDGMENT AND BEHAVIOUR:

A descriptive and prescriptive analysis

Jeremy David Cone, Ph.D.

Cornell University 2012

Researchers have had a long-standing interest in the relative contributions of implicit and explicit processes to judgments and behaviours. Although conscious deliberation has often served as the “gold standard” for normatively correct decisions, researchers are increasingly beginning to focus their attention on the relative strengths of implicit processes, pointing to situations in which, it is claimed, they can sometimes exceed rational deliberation. In this work, I explore both the descriptive question of how individuals come to trust the dictates of an implicit or explicit process when the two pull in opposing directions, as well as the prescriptive question of when implicit and explicit processes can serve as reliable guides for sound judgments and adaptive behaviours.

Three lines of work examine these questions. In the first line (Studies 1-3), I ask how people come to trust intuition or reason in the context of situations in which they experience a conflict between them. I find consistent support for what I call the *state cuing hypothesis*—that is, the notion that people attempt to resolve such conflicts by looking for cues in their current mental states that help them decide what to do. Specifically, I find that when the features of one’s mental state matches the features people associate with intuitive or rational thought, this serves to make one source of input “feel right,” causing it to exert a greater influence on the decision.

In the second line of work (Studies 4-8), I assess the extent to which intuitions can successfully capture relevant aspects of one's prior experiences with decision objects, thus allowing them to serve as a reliable source of information when choosing among them. However, across 4 replication attempts, I find that intuitive judgments are perhaps not as reliable as past research would suggest.

Finally, in the third line of work (Studies 8-12), I assess whether implicit attitudes can be quickly revised in light of recent, relevant experiences. I find consistent support across the studies that implicit attitudes can indeed be rapidly revised in light of new, countervailing information, thus suggesting that implicit attitudes can serve as a reliable guide for behaviour even in situations in which recent experience is inconsistent with past learning.

Taken together, this research attempts to advance our understanding both of the factors that can influence the extent to which people believe that implicit and explicit processes are valuable and reliable, as well as the actual contribution of each type of process to normative and satisfying judgments and behaviours.

BIOGRAPHICAL SKETCH

Jeremy Cone was born in Richmond Hill, Ontario, Canada. He came to Cornell's Social Psychology program in something of a roundabout fashion, first receiving a degree in Computer Science from the University of Waterloo in 2005. After spending a year living the real-life version of Office Space, he realized very quickly that he was especially ill-suited to a life in the cubicles, returning to the University of Waterloo in the Fall of 2006 to pursue a second degree in Psychology and then ultimately finding his way to Cornell in Fall 2007. He begins his re-entry into the "real world" at Williams College in Williamstown, MA.

ACKNOWLEDGEMENTS

To call this “my” dissertation seems particularly dishonest in light of the amazing collection of people that have helped to get me to this particular point in July 2012, both directly, through their superb guidance and expertise at each step along the way, and indirectly, through their amazing level of support and encouragement throughout the process.

I want to thank, first and foremost, my fiancée, Cait, who has been kind enough to tolerate each and every discussion over the years of the precise nuances that distinguish Gawronski’s Associative-Propositional Evaluation model from Wilson’s dual attitudes model (and various other minutia) as though it was the most interesting she had ever had; every $p > .05$ and the concomitant instantaneous sharp increase in negative affect and self-doubt; and every day over the last five years that we’ve had to spend apart so that I could live the dream of pursuing a doctorate in psychology at Cornell. Thanks, sunshine, for being an unending source of support, love, thoughtfulness, and patience. You’ve been there to celebrate every success and to remind me of What’s Really Important after every failure, even from five hundred kilometres away. I can’t imagine this document being anything but a series of blank pages without you. I absolutely can’t wait to marry you—and, of course, for zero kilometers.

My primary advisor, Tom Gilovich, has been the most consistently inspiring and unendingly supportive mentor that I could have possibly hoped for. In job interviews, whenever I’m asked about my teaching philosophy, I’ve been known to launch into a brief monologue about the value of communicating not just what psychology has to teach us, but also why anyone would ever want to study it in the first place. Tom represents the best role model that I know for exactly that kind of inspiration. Every time we meet, I’m a witness to what it is to relish in seemingly every aspect of being an academic, and I can’t help but relish in them too. Even when we’re discussing failed studies or other ostensible dead ends, I always feel, at the end of our meetings, as though the “Way Forward” is readily apparent—and should have been all along, if only I could have seen the problem with his vision and approached it with his clarity and purpose. I’m confident that I could spend another five years under his guidance and continue to learn and grow as much as I have in the previous five.

To say that I work with Melissa Ferguson is something of a misnomer, because it never feels like work at all. Our meetings—and, more often, our flurried e-mail exchanges that quickly balloon out of control—have represented some of the most intellectually exciting moments of my graduate career. Whether it’s considering the exact meaning of the word “change”, or discussing the precise nuances that distinguish Gawronski’s Associative-Propositional Evaluation model from Wilson’s dual attitudes model, I always come away from our conversations with at least three new things to consider, and at least four new articles to read. She represents much of what I aspire to be as I transition into the next phase of my career—as a researcher and colleague, as a teacher and mentor, and as a friend.

David Dunning has taught me, in scientific research, the value of patience, thoughtful analysis, and an especially well-constructed play on words. Some of my most memorable “teachable moments” of the last five years occurred after spending several hours toiling away at a data set from one of our studies—trying to uncover some hint as to why a study result had turned out as it had—only to discover an e-mail from Dave with his proposal for the precise set of next steps, just as a doctor would effortlessly prescribe a medication for the easiest and most common of ailments. Never once shying away from giving any problem he encounters a really “good think” (by which I mean several hours—often days—of intense, concentrated effort), he represents for me the joy of living a life of the mind in its purest, most unadulterated form.

To my breadth member, Shimon Edelman, I owe a tremendous debt of gratitude for always keeping me honest, and for always encouraging me to seriously question even the most heavily entrenched assumptions and conventional wisdom of our discipline. The simple act of imagining, “what Shimon might say” to the claims I’ve made throughout this dissertation has undoubtedly made it a much better piece of work.

I’ve been tremendously lucky to have the support, kindness, generosity, encouragement, and wisdom of such a wonderful group of people. It seems wrong that such an amazing and intellectually fulfilling experience should come to an end—that I should be leaving now to head somewhere else. I can only hope to experience—and to have some hand in recreating—the kind of magic and excitement that Cornell’s Social Psychology program embodies in every place I go from here. I certainly plan to try.

TABLE OF CONTENTS

Biographical Sketch	i
Acknowledgments	iii
Table of Contents	iv
List of Figures	v
List of Tables	vi
Chapter I. Introduction	1
Part I: Description • Which Process Do People Rely Upon?	
Chapter II. State Cuing: Mental states cue choice strategies	15
Part II: Prescription • When Can Each Process Be Relied Upon?	
Chapter III. On the robustness of intuitive impulses	44
Chapter IV. Changing your implicit mind: On the rapid formation and revision of implicit attitudes	69
Chapter V. General Discussion	113

LIST OF FIGURES

Figure 1. Participants' responses on the urn and apartment choices in Study 1.	29
Figure 2. Participants' responses to the urn and apartment choices in Study 2.	34
Figure 3. Screenshot of the stock ticker program used in Studies 4-7.	54
Figure 4. The last two lines of participants' purported data file for a participant initially assigned to prefer <i>Xanthie</i>	80
Figure 5. Implicit Responses (D-Scores) towards Quan/Xanthie in Study 8.	83
Figure 6. The last four lines of participants' purported data file, in this case for those initially given feedback that they were an Underestimator (in Study 9a).	89
Figure 7. Participants' implicit preferences for novel group 1 over novel group 2 before (time 1) and after (time 2) learning of the "mistake" in their computer feedback. .	92
Figure 8. A wug.....	94
Figure 9. The two-way interaction between time and round in Study 11. This effect was unmoderated by amount of exposure to <i>wugs</i>	97
Figure 10. Participants' proportion-pleasant judgments at time 1 and time 2 in Study 12.	100
Figure 11. Participants' proportion-pleasant judgments towards Bob before (time 1) and after (time 2) learning either that he was convicted of molesting children (experimental) or that he bought a soda (control) in Study 12.	106

LIST OF TABLES

Table 1. Intuitive judgments of stocks as a function of actual performance in Study 4	56
Table 2. Intuitive judgments of stocks as a function of actual performance in study 5	61
Table 3. Intuitive judgments of stocks as a function of actual performance in Study 6	63
Table 4. Intuitive judgments of stocks as a function of actual performance in Study 7	65

CHAPTER I. Introduction

I.

By any reasonable measure, Timothy Cook had, by the beginning of 1998, already achieved an impressive amount of success. Cook was the Vice President for Corporate Materials at Compaq—a company that, in the booming IT industry of the 1990s, had become an exceptionally profitable manufacturer of desktop computers. Its flagship product, the Compaq Presario, was a household name—one of the first personal computers available for less than a thousand dollars, making it a more affordable alternative to higher-end products of that era.

Early that year, an executive head-hunter approached Cook with an opportunity. At the time, Apple—one of Compaq’s main rivals in the industry—was in the midst of a dramatic re-shuffling of its corporate leadership. Just six months earlier, the company had ousted its CEO, Gil Amelio, and replaced him, in a strange turn of events, with one of the men that had founded the company in the 1970s but that had, just over a decade earlier, also found himself ousted by the board of directors—Steve Jobs. Back at the helm of the company he created, Jobs was looking to reinvent the Apple brand. He had a strong vision about how to put the technology company back on top again. And he wanted Cook to be one of the people to help him do it.

Faced with a choice between staying in his current position at Compaq and following Steve Jobs in his vision of reinventing the company he founded, Cook, being an engineer by training, knew the “right” way to make the decision.

When he enumerated all of the costs and benefits of each potential course of action, the decision could not have been more obvious. The Apple of today is such a powerhouse of the IT world—with dizzying numbers of fans lining up for hours to buy their latest products on the day of their release—that we largely forget that the Apple of the 1990s was, by all accounts, a company in complete disarray. In 1998, while Cook was attempting to make his decision, the handheld devices that have come to be most strongly associated with the Apple brand—the iPod, iPhone, iPad, even the iMac—were still vague ideas on whiteboards. It was, according to Time Magazine, “one of the worst managed companies in the industry.” The company had suffered a number of high-profile financial losses and product flops. Its stocks had sat at record lows for the previous three years. It was quickly becoming a pariah of the industry, widely considered to be a failing brand just barely on life support. When Michael Dell, the founder of another well-known technology giant of that era, was asked at ITxpo '97 for his thoughts about what could be done to make Apple profitable again, he famously said, “I’d shut it down and give the money back to the shareholders.”

By contrast, the Compaq of 1998 could not have been more vibrant or its future seemingly brighter. By the time that Cook had been approached by Jobs, his current employer had exceeded even IBM as the largest manufacturer of personal computers in the world. The company had been turning record profits in quarter after quarter for the previous several years. Forbes had just crowned

it with the designation of Company of the Year. Others had called it the “king of the business world.”

But something about changing companies just felt right. For reasons that escape him to this day, five minutes into his meeting with Steve Jobs, Timothy knew that he wanted to work at Apple. Though his friends and colleagues told him that he would be a fool to take the job, his intuition told him that Steve Jobs was right and that his friends and colleagues—not to mention the actuarial tables, and the market analysts, and the costs and benefits analyses, and Michael Dell—were wrong. That was how, in March of 1998, Cook became the Senior Vice President of Worldwide Operations at Apple, working alongside Steve Jobs in one of the most dramatic corporate turnarounds in recent memory.

Today, Cook has, in the wake of Steve Jobs’ passing, stepped into the role of CEO of the world’s most recognizable tech brand, and, with his acquisition of one million of Apple’s stock units as part of the compensation package he received at the beginning of this year (valued at over \$400 million), he is now the world’s highest-paid CEO.

Compaq, however, has fared considerably less well. Just two years after Cook’s departure, as a result of a high-profile merger with Digital Equipment Corporation that turned sour and the realities of a market that was steadily becoming less interested in personal computers in favour of the handheld devices that Apple would come to represent, the company began to falter. Not long after being named Company of the Year, Compaq was, rather suddenly and

rather jarringly, being described by market observers as “a shaken company with an uncertain future.” In the final two-and-a-half years of the 1990s—while the ‘i’ prefix that has come to symbolize Apple’s innovative product line was achieving its first market presence with the release of the iMac in 1998—Compaq lost more money than it had in the previous sixteen years combined.

Twelve years later, describing a decision that, with the benefit of hindsight, we can now recognize as the single best decision of his career despite every rational consideration that would have been available to him at the time, Cook said, “there are times in our lives when the careful consideration of costs and benefits just doesn't seem like the right way to make a decision. There are times in all of our lives when a reliance on gut or intuition just seems more appropriate—when a particular course of action just feels right.”

II.

Stories of the remarkably insightful intuitions of highly successful people possess a quality of mystery and intrigue. We delight in learning of CEOs that somehow “knew” what others did not—situations in which a lone individual goes against the advice of everyone around her and every rational consideration available to her but is nonetheless later vindicated for “trusting her gut” when she somehow knew it was right. Many read such stories in the hopes that they too can somehow, in Gladwell’s (2005) terms, “think without thinking” and arrive at stunningly successful choices and decisions. In this way, intuitive judgment possesses many of the qualities that Chabris and Simons (2010) have called “the

illusion of potential”—the cognitive illusion that we can rather effortlessly acquire new skills or abilities. If only we could learn to harness the power of our intuitions and better learn to “listen to our hearts,” we too could know that leaving the “king of the business world” for “one of the worst managed companies in the industry” would be one of the best decisions of our careers. Indeed, if the recent success of a number of popular treatments of intuition is any indication (Gladwell, 2005; Gigerenzer, 2007; Myers, 2002), people seem quite willing to accept the idea that intuitions can be a successful, stunningly prescient source of information that can aid decision-making if people are willing to place their trust in them.

Outside of popular culture, however, scholars have been relatively more reluctant to grant any kind of power or mysterious, ineffable insight to intuitive judgment or other types of non-conscious processes. Sound rational analyses have long been considered the gold standard of high-quality decision-making, and it is generally thought that individuals ignore such rational considerations at their peril. At the same time, passion and emotion have, at least in scholarly circles, long been considered a source of bias and error that must be silenced in favour of dispassionate logical analyses. Psychologists have been strongly influenced by such thinking. Indeed, in summarizing researchers’ views of intuitive judgment across many decades of research, Hogarth (2010) has observed that it “has not been held in high esteem by researchers in most of the last 60 years” (p. 339).

The reasons for researchers' relative skepticism about the quality of intuitive judgment are not especially hard to discern. Early examinations of individuals' judgments had demonstrated that even a simple actuarial formula that took account of only a small subset of potential factors that could influence a decision could invariably outperform participants' idiosyncratic clinical predictions (Meehl, 1954; Dawes, Faust & Meehl, 1989). Moreover, even when a statistical formula resulted in systematic prediction errors, participants' clinical judgments were nonetheless inevitably found to be even more error-prone. Indeed, in decades of research, no situation had been uncovered in which participants' unaided judgments consistently exceeded the performance of a relevant rational analysis (Dawes, Faust, & Meehl, 1989). These early demonstrations had thus strongly implied that rational deliberation was an invariably better alternative to whatever other idiosyncratic knowledge participants may have brought to a decision, and that individuals' snap judgments might best be ignored.

Another reason for researchers' skepticism of the power of intuitive judgment comes from the highly influential heuristics and biases program of research on judgment under uncertainty pioneered by Kahneman and Tversky in the 1970s (Tversky & Kahneman, 1974; see also Kahneman, 2011). The success of this program of research in identifying a host of situations in which individuals had especially compelling intuitions that were normatively inferior to a reasoned analysis serves as a strong testament to the value of clear, rational thought. For

example, in one of Tversky & Kahneman's (1983) most famous demonstrations, participants were asked to consider the following problem:

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

Which is more probable?

1. Linda is a bank teller.
2. Linda is a bank teller and is active in the feminist movement.

Participants' intuitions strongly compel them to believe that the second option is more probable (85%) even though this is a logically impossible outcome by virtue of it being a subset of the first.

Kahneman and Tversky (1982) were careful to point out that their focus on the biases and errors inherent in people's intuitive judgments under uncertainty was ultimately in the service of understanding the mechanisms by which such judgments operate. In the same way that visual illusions can tell us something important about the operation of an otherwise exceedingly adaptive and well-functioning visual system, they said, so too can the times when intuitions lead to compelling yet normatively unsound conclusions tell us something important about the nature of people's judgments and decision-making. Indeed, they acknowledged that this exclusive focus on the situations in which intuitions falter could give a false impression about the relative quality of such judgments, suggesting that heuristics serve as reliable guides of behaviour much of the time but could nonetheless lead to error under the well-defined circumstances captured in their various experimental dilemmas. However, the

large accumulation of results of the sort described above has, despite Kahneman and Tversky's caution, nonetheless had an important influence on researchers' suspicions about the quality of intuitive insight (but see Gigerenzer et al., 1999).

Many lines of seminal research in judgment and decision-making and a long tradition in Western thought have thus cast intuition in a relatively poor light. Interestingly, however, a number of researchers in recent years have sought to shift the field's attention away from the shortcomings of intuitions and toward the relative utility of non-conscious processes. This has resulted in the identification of a number of situations in which people's gut feelings seem to give rise to adaptive decisions and has brought to light a wide range of stunningly complex and adaptive judgments that appear to be the product of trusting the "gut" and "listening to the heart." Indeed, investigators have now begun cataloging a number of instances in which intuitions appear to lead individuals not only to normatively correct or subjectively satisfying decisions, but even, it is claimed, to better outcomes than those obtained by individuals relying on more deliberative judgment (Betsch, Plessner, Schweiren, & Gutig, 2001; Dijksterhuis, 2004; Dijksterhuis, Bos, Nordgren, & van Baaren, 2006; Dijksterhuis & Nordgren, 2006; Klein, 1999; Wilson & Schooler, 1991). Outside of the judgment literature, non-conscious processes have similarly been implicated in a number of highly complex and adaptive behaviours (see, e.g., Bargh & Ferguson, 2000; Chartrand & Bargh, 2002; Fitzsimons & Bargh, 2004; Hassin, Bargh, & Zimmerman, 2009).

In this way, researchers' recent attention has begun to shift from the shortcomings of implicit processes to the relative strengths and weaknesses of each type of process. However, this research is still very much in its nascent stages and we still know relatively little about when and why intuitive judgments (or the behaviours that are the product of other non-conscious processes) can serve as a basis for sound decisions. In this dissertation, my goal is thus to advance our understanding of how implicit and explicit processes interact to influence judgment and behaviour, as well as to better our understanding of the normative question of when implicit processes are valid and reliable and when they are more likely to be a source of bias and error that should be discounted in favour of explicit processes.

My discussion will proceed in two parts. In the first part, I will explore the descriptive question of how people come to choose a decision strategy—that is, whether they ultimately choose to rely on one process or the other when making a decision, particularly in situations in which the two processes pull in two different directions. In the second part, I will explore the prescriptive question of when and why implicit or explicit processes may be more reliable and adaptive in a given situation. This is naturally a very large and complex question and there will be a large host of factors that will ultimately come to influence whether the products of implicit processes are veridical or not. My discussion of this question will thus be somewhat circumscribed, focusing on a few key assumptions about

the nature of implicit and explicit processes and exploring the implications of these assumptions for the quality of judgment and behaviour.

VI. Outline of subsequent chapters

Chapter II

In Chapter II, I explore the descriptive question of how people come to choose a decision strategy in the context of judgment, exploring, in particular, situations in which people face a conflict between the dictates of their gut reaction to a decision problem and those of a more rational, deliberative analysis. What is it that leads people to trust their intuitions in some cases—such as Timothy Cook’s surprising decision to leave the “king of the business world” for “one of the worst managed companies in the industry”—and to trust reasoned analysis in others?

In particular, I empirically test what I call the *state cuing hypothesis*—that is, the prediction that what makes intuition or reason “feel right” in a particular situation is the extent to which people can turn to cues in their current mental states that closely match the properties of intuitive or rational thought. When people are thinking quickly—a property that closely matches the speed and spontaneity of intuitive thought—people are more inclined to trust their intuitions; when people experience mental states that encourage more deliberative analysis, such as a feeling of meta-cognitive difficulty in processing information about a problem, they are more inclined to trust rational analyses.

In this chapter, I refrain from drawing any conclusions about the relative quality of people’s decision-making made on the basis of intuition or reason, and instead focus on the predictors of the particular choice strategy that people adopt in a given situation.

Chapter III

In the second part of the dissertation (Chapters III and IV), I turn my attention to a prescriptive analysis of the circumstances in which implicit and explicit processes can serve as the basis for sound decision-making. In Chapter III, I explore this question in the context of intuitive judgment, examining the factors that have been posited to influence the quality of intuitive insights. My analysis in this chapter hones in on one key assumption that has been made about the processes that govern the development of gut feelings—that is, that they provide a summary evaluation of the breadth of one’s past experiences with a decision object. Researchers have argued that this ability to capture all of one’s prior experience can give intuitions an upper hand in situations in which this past experience can usefully be applied to one’s later decisions (e.g., Betsch et al., 2001).

However, even if intuitions *can* capture such information, an important theoretical question with respect to their usefulness in everyday judgment is whether the processes that govern the formation of intuitive impulses can discount aspects of one’s past experience that are based on fallacious information or are irrelevant for later judgments. That is, their usefulness is

rather limited if they do not possess a necessary validity filter and instead promiscuously incorporate all aspects of one's previous experiences into a summary evaluation irrespective of their truth value. I sought to submit this question to empirical test by training participants' intuitions in a well-known experimental paradigm, introducing a source of biased and irrelevant information for some participants and not for others. The main question of interest in the studies I report here was thus whether participants' intuitions differed under these different learning conditions or if intuitive judgments remained robust to their influence.

Chapter IV

In Chapter IV, I broaden my prescriptive analysis by turning my attention to the attitudes literature—another area of research in which the interaction between implicit and explicit processes plays a prominent role. My focus in this chapter is on an assumption that has been made about the nature of implicit attitudes that has important implications for their reliability—that is, that they are relatively slow to develop and relatively heavily-entrenched once established. This assumption implies that implicit attitudes are rather insensitive to recent yet highly relevant experiences, making them a source of bias and error in situations in which these recent experiences are necessary for adaptive behaviour. Such a situation may arise, for example, when one learns that one's romantic partner of several years has been having an affair with one's best friend. In this situation, recent experience is both: (a) (presumably) inconsistent

with prior experiences, and (b) highly relevant for future behaviour. Yet, because implicit processes are rather insensitive to such information, they may serve as a source of maladaptive behaviour towards one's partner.

My focus in chapter IV is on more thoroughly examining the theoretical underpinnings of this assumption about the nature of implicit attitudes, exploring whether they must necessarily be a source of bias in such situations. Based on theoretical conceptions of associative memory, I posit that there are, in fact, a number of situations in which implicit attitudes can indeed quickly incorporate recently-learned relevant information, and provide empirical evidence to support this assertion.

**PART I. Description • Which Process
Do People Rely On?**

CHAPTER II. State Cuing: Mental States Cue Choice Strategies

Timothy Cook told the audience at his commencement address at Auburn University in 2010 that if he had not cast aside the advice of the people that knew him best—if he had trusted his rational analyses of his prospects at Apple—he would not have been giving a speech that day and very few people would have known who he was. Moreover, when Cook sided with his intuition and with Steve Jobs, he was going against everything that his training as an engineer had taught him to do in such situations. He had done the appropriate cost-benefit analysis and yet somehow found it lacking. What exactly was it that made Cook trust his intuitions that day? How was it that this major career decision came to be, in his words, one of the “times in his life when the careful consideration of costs and benefits just didn’t seem like the right way to make the decision”?

As Cook’s situation in March of 1998 highlights, such decisions often carry a great deal of weight. There are times, like Cook’s, in which stunningly complex and adaptive judgments appear to be the product of trusting the “gut” and “listening to the heart”—a point that is substantiated by many instances in the judgment literature in which intuitions appear to lead individuals to highly adaptive and satisfying choices (Betsch, Plessner, Schweiren, & Gutig, 2001;

Dijksterhuis, 2004; Dijksterhuis, Bos, Nordgren, & van Baaren, 2006; Dijksterhuis & Nordgren, 2006; Klein, 1999; Wilson & Schooler, 1991).

But as helpful as intuitions may sometimes be, it is nonetheless easy to cite countless instances of times when they have led to faulty judgments, erroneous beliefs, and non-optimal choices (Baron, 1988; Gilovich, 1991; Gilovich, Griffin, & Kahneman, 2002; Hall, 1980; Kahneman, 2011; Kahneman, Slovic, & Tversky, 1982). That is, it is easy to cite situations in which intuitions should have been discounted in favor of a more deliberative, rational analysis. Indeed, even the most ardent supporters of intuitive decision-making are careful to point out circumstances in which it is better to ignore or suppress gut feelings (e.g. Dijksterhuis, Bos, Nordgren, & van Baaren, 2006).

Thus, to effectively navigate a complex world in which gut feelings sometimes aid and sometimes hinder judgment and decision making, individuals must develop intuitions about their intuitions – that is, beliefs about when their gut-feelings can be trusted and relied upon for a sound decision, and when it might be best to ignore such feelings and defer to a rational analysis. Should people pay attention to a nagging feeling that tells them to dump a particular stock, or is it better to focus on a long-term projection indicating that it would be better to hold onto it? Should a father trust his instincts about his daughter's new beau, or is that likely to fall into the category of being overprotective?

Despite a large and growing literature on the relative value of intuitive judgment, relatively little is currently known about how people come to trust

one source of input into the decision-making process over another. Are there factors that can make intuitive choices seem more compelling? In Chapter II, my goal is to examine this question in more detail, exploring whether there are predictable influences that can cause one process to pull harder, ultimately encouraging people to trust the dictates of one process over the other.

II. Dual process models of judgment

Researchers have generally conceptualized intuition and reason as the product of two distinct modes of thought. Although different theorists have put forward models that differ somewhat in their concrete details, they generally agree on the properties that characterize each mode of thought (Epstein, 1991; Evans, 2004; Kahneman, 2003, 2011; Sloman, 1996; Strack & Deutsch, 2004; but see Reyna & Brainerd, 2011, for a rather different dual process model). One mode—governed by what has been variously labeled the experiential system (Epstein, 1991), associative system (Sloman, 1996), or simply System 1 (Kahneman & Frederick, 2002; Kahneman, 2011; Stanovich & West, 2000)—is said to be automatic, fast, associative, affect-laden, non-conscious, and effortless. The other—governed by what has been labeled the rational system, rule-based system, or System 2—is said to be deliberate, slow, rule-based, conscious, and effortful. These different accounts also share the key assumption that these two modes of thought can operate in parallel and sometimes produce conflicting outputs, giving rise to the quintessential intuition/reason conflict that

people sometimes face. As Sloman (1996, p.19) put it, “The fact that people are pulled in two directions at once suggests two forces pulling.”

Of course, although these models represent the dominant perspective in research in judgment (and many other areas of social psychology), they are not without controversy (for critiques, see Keren & Schul, 2009; Kruglanski & Thompson, 1999; Gigerenzer & Regier, 1996). Dual process models have been strongly challenged on a number of grounds, including the extent to which the properties that distinguish the two systems actually represent binary dichotomies rather than continuous variables (e.g., “automatic vs. controlled”); whether the similarities among the models I highlighted above are more apparent than real; and the extent to which the two systems are isolable (for a detailed discussion of these issues, see Keren & Schul, 2009; Kruglanski & Thompson, 1999; Gigerenzer & Regier, 1996).

Whether these two modes of thought are best characterized as the product of one system, two systems, or more, however, what is relatively less controversial is that individuals have very different subjective experiences of the two modes of thought. Suppose, for example, that you learn that it will be 20°C tomorrow. Will you need a sweater? For those who have extensive experience with the Celsius temperature scale, the answer comes to mind spontaneously and with very little thought. For those who have relatively little experience with this scale, however, answering such a question involves relatively more cognitive effort, likely requiring the application of a mathematical formula to convert the

temperature from the Celsius scale to the Fahrenheit scale ($F = \frac{9}{5}C + 32$)—a value that can then be used to intuit a response. This subjective difference is the one that I will hone in on in this chapter, and I will remain relatively agnostic about the origins of these subjective differences.

III. Going with the head or the gut

Like all complex problems, there are likely to be a number of important factors that influence which of the two forces ultimately exerts a stronger influence on decision-making. For one thing, it seems odd to choose a romantic partner by trying to carefully weigh various “objective” pros and cons or by attempting to calculate the precise probability, based on each potential partner’s life history, that one will end up blissfully happy, content, disgruntled, or divorced. Intuitively, such decisions seem less amenable to a rational analysis than, say, choosing a car stereo or deciding which of several models of iPod to purchase—decisions in which the attributes of each alternative can be delineated and the quality of each alternative precisely evaluated. And even if one does engage in this sort of rational analysis to choose the best partner and that analysis conflicts with a gut feeling, most people would argue that in a decision like this it is better to go with the gut feeling. But why? What makes choosing a mate an intuitive domain and choosing a stereo a rational domain?

IV. Task & State Cuing

In previous work, we (Inbar, Cone & Gilovich, 2010) examined the factors that make intuitions seem especially diagnostic in some domains and rational

analysis more appropriate in others. We found that the tendency to favour intuition or reason can be predicted from the features of the decision context or the object being chosen. More specifically, we found that people appear to be cued one way or the other by the extent to which the features of the decision or decision context resemble features people associate with intuition or reason. For example, rational judgment is thought to be objective (in that it involves the application of well-specified rules and criteria) whereas intuition is seen as subjective (in that it entails personal, idiosyncratic preferences). When evaluating whether one chose well or poorly is similarly subjective (e.g., one's satisfaction with an entrée or a romantic partner) people tend to favour their intuitions when making a decision; when the outcome can be objectively evaluated (e.g., the success of a military campaign or the return on an investment), people tend to favour reason. Similarly, because intuition is experienced as a holistic, all-or-none flash of inspiration, whereas reason is a slow, step-by-step process, people tend to follow their rational analysis for decisions that tend to proceed in a series of sequential stages and their intuition for those that happen "all at once."

Individuals thus appear to be cued by aspects of the decision context such that one mode of thought simply "feels right" and they side with that mode when resolving intuition/conflicts—a process that we have called *task cuing*. In the current studies, I explore whether people are similarly cued by aspects of their internal state when deciding what to do – a process I call *state cuing*. I

propose that people look for clues for how to resolve intuition-reason conflicts wherever they can find them: in much the same way that people look outward to features of the decision context to decide which source of input to trust, they also look inward to features of their current mental state to discern which course of action “feels right.” For example, because rational analysis is more effortful and deliberate than intuitive judgment, the sense of engaging in effortful processing should lead individuals to side with their rational analyses. Conversely, because intuitive judgment is rapid and efficient whereas rational analysis is slow and laborious, the sense that one is thinking at a crisp, rapid pace should foster a preference for intuition.

This proposal fits with a considerable amount of recent work on the impact of meta-cognitive assessments on judgment and choice, and with work showing that there is considerable “cross talk” between separate mental representations and processes, with inputs or computations in one area having an impact on assessments and outputs in other, seemingly unrelated areas. An incidental physical sensation of weight, for example, can influence the significance, or weight, that people attach to stimuli (Jostmann, Lakens, & Schubert, 2009); the sensation of physical warmth can make a target person seem warmer (Williams & Bargh, 2008) and a friend seem closer to the self (Ijzerman & Semin, 2009); and being reminded of an episode of rejection can make the ambient temperature seem colder (Zhong & Liljenquist, 2006).

More generally, a “fit” between one’s current mental state and a to-be-evaluated stimulus has been shown in numerous circumstances to lead people to infer that a choice is right or a proposition is valid (Cesario & Higgins, 2008; Higgins, 2005; Higgins & Scholer, 2009; Johnson & Tversky, 1983; Lee & Aaker, 2004; Risen & Critcher, 2011; Schwarz, Sanna, Skurnik, & Yoon, 2007). For example, Pham & Avnet (2004) found that when people are in promotion-focused mindsets (that is, those in which they eagerly seek out successes and rewards), they are relatively more likely to make their choices based on affective responses to alternatives, whereas when they are in a prevention-focused mindset (that is, those in which they vigilantly avoid failures and punishment), they tend to favour making decisions on the basis of more substantive, reason-based information. The researchers argue that participants’ differential sensitivity to these two types of information are driven, in part, by the fact that prevention-focused mindsets encourage increased vigilance and risk-aversion, which ultimately serve to increase reliance on analytical processes—an effect that, from the perspective of the state cuing hypothesis, makes rational judgments “feel right.”

In preliminary tests attempting to assess the state cuing hypothesis more directly, Inbar (2008) found that when participants are made to feel accountable—a mindset in which rational justification of one’s choices is deemed important and necessary—participants are cued to trust their rational analyses. Similarly, when a hypothetical decision is described with a number of concrete,

vivid details—properties that best characterize intuitive thought relative to the more abstract nature of deliberative analysis—participants become relatively more inclined to side with their gut.

Although these initial demonstrations are consistent with state cuing, the evidence for this hypothesis rests not on any one study, each of which may have any number of alternative explanations, but rather on the consistent picture painted by a diverse collection of manipulations that employ many different types of cues and many different types of decisions in which intuition and reason pull in opposing directions. The current studies thus sought to provide direct tests of the state cuing hypothesis, inducing states that matched quintessential features of rational and intuitive thought, and assessing participants' tendencies to choose in line with their intuitive assessment or a more deliberative one.

V. The Current Studies

The current studies assessed state cuing by employing both hypothetical choices and actual decisions with real monetary consequences. In each study, participants were first placed into a particular mindset that closely matched one of the widely-accepted features associated with intuitive and rational judgment. They were then asked to make a decision that I thought would induce an intuition/reason conflict and I examined whether they chose to follow their head or their gut. In Study 1, I manipulated perceptual fluency to induce a meta-cognitive sense of ease or difficulty of thought and examined how it influenced participants' responses in the ratio-bias paradigm (Denes-Raj & Epstein, 1994;

Epstein, 1991) and in a hypothetical apartment choice. In Study 2, I examined whether a meta-cognitive sense of thinking rapidly would make participants favour their (rapid) intuitive impulses. Finally, in Study 3, I examined the effect of the physical “closeness” of a decision on preferences for intuition over reason. Because people feel that that intuitions come from deep within themselves, whereas reason is something that exists outside the self (because it relies on learned rules and much of it can be done on a blackboard, on a pad of paper, or with a calculator), decisions that are presented “in close” to the body might elicit more intuitive responses than those that are presented at a distance.

Study 1: Meta-Cognitive Fluency

Rational analysis is typically characterized as an effortful, deliberate process whereas intuition is seen as spontaneous and effortless. The state cuing hypothesis therefore entails that a feeling of engaging in more effortful or deliberative processing should serve as a cue to the individual that a rational analysis is a good fit to the task at hand, which should lead the individual to resolve an intuition-reason conflict in favor of reason. A large literature indicates that processing (dis)fluency leads to exactly such a state (Alter & Oppenheimer, 2009). That is, processing disfluent stimuli leads to a meta-cognitive sense of effort and difficulty—a feeling that, according to the state cuing hypothesis, should lead people to favor their reasoned analysis.

Indeed, Alter, Oppenheimer, Epley and Eyre (2007) found that meta-cognitive disfluency appears to prompt individuals to engage their rational

faculties and to be more careful and deliberative in their analysis of a problem or decision. In one study, participants completed the Cognitive Reflection Task (CRT) – a test which consists of 3 items, each of which has an intuitively compelling but incorrect answer that must be overcome with deliberate re-consideration. When the CRT was presented to participants in a hard-to-read font, they were less likely to give the intuitively appealing but incorrect answers. In another study, when a persuasive advertisement was difficult to read, participants relied more on systematic processing (that is, the central route in the Elaboration Likelihood Model—Petty & Cacioppo, 1986; Petty & Wegener, 1999) than when the advertisement was in an easy-to-read font.

Oppenheimer and colleagues' research indicates that the meta-cognitive sense of disfluency leads people to engage in more deliberative thought about a problem. The state cuing hypothesis takes this a step further: Disfluency should not only lead people to engage in more deliberative thought but it should also lead them to *trust* the outputs of their rational analysis more. That is, even when the bulk of the processing is done and one is nevertheless torn between an intuitive option and a rational option, a meta-cognitive sense of disfluency should lead people to find the rational option more appealing.

To test this idea, I had participants make choices in a version of the ratio-bias paradigm, a choice dilemma well-known for creating an intuition/reason conflict in many participants. Notably, in this paradigm, participants know the normatively correct choice—they can readily identify and articulate which option

is rationally superior. Nonetheless, they often end up favouring the inferior option. Thus, unlike the Cognitive Reflection Task in which greater deliberation makes it more likely that participants will arrive at the correct answer, in the ratio-bias paradigm participants tend to know the correct answer but nonetheless have a hard time selecting it (e.g., Kirkpatrick & Epstein, 1992, Study 3). Their choices, then, depend on whether they are inclined to favour one input—the rational or the intuitive—over the other.

Method

Fifty-three Cornell undergraduates completed an online questionnaire as part of their participation in an unrelated experiment. Participants were randomly assigned to either a *fluent* condition, in which the text of the survey was presented in a standard, easy-to-read font (Calibri, 11pt; Sample: Calibri), or a *disfluent* condition, in which the text was presented in a difficult-to-read font (Haettenschweiler, italicized, 11 pt; sample: *Haettenschweiler*). Previous research has used similar font manipulations to induce a state of perceptual disfluency (see Alter & Oppenheimer, 2009).

Participants first read a version of the ratio-bias paradigm (Epstein, 1991). In this paradigm, participants are presented with two urns, each containing two different colours of tokens (usually marbles or jellybeans)—a winning colour and a losing colour. They are told that a token will be drawn randomly from the urn of their choosing and, if a winning token is drawn, they will win a small sum of money. One of the urns contains 1 winner and 9 losers. The other urn always

contains 100 tokens, but the precise number of winners varies, typically between 5 and 9. The key feature of the larger urn is that it always contains a higher absolute number of winners but a lower percentage. Research using this paradigm has established that many participants experience a conflict between the rational choice (the small urn with better odds of winning) and the intuitive choice (the large urn with a greater number of winners) (Epstein, 1991).

In my study, participants completed a computerized version of the ratio-bias paradigm. The text at the top of the screen read: “You are going to draw from one of the two urns depicted below. If you draw a red marble, you will win \$20. Please decide which urn you would like to draw from. Click on your choice.” Beneath the description was a two-dimensional graphical depiction of two urns. The small urn contained 9 white (losing) marbles and 1 red (winning) marble. The large urn contained 91 white (losing) marbles and 9 red (winning) marbles. Below each urn, the text indicated the probability (10% and 9%) of drawing a winning marble, each displayed in the same fluent or disfluent font as the instructions, depending on condition.

After participants clicked on one of the two urns, they read a hypothetical scenario in which they were asked to imagine they were deciding which of two apartments to rent: “Imagine that you are choosing an apartment between two options. Imagine that rational deliberation has led you to favor Apartment A, but you feel a strong intuitive draw toward Apartment B. Which apartment would

you choose?” Participants clicked on options labeled either “Apt A” or “Apt B”, presented in the same font as the instructions.

Results

To assess the extent to which participants were inclined to follow the dictates of a rational analysis or go with a gut feeling, I performed a 2 (fluency: fluent or disfluent) × 2 (choice dilemma: apartment or ratio-bias paradigm) mixed-model analysis of variance, with fluency as a between-subjects factor and choice dilemma as a within-subjects factor. This analysis yielded the predicted significant main effect of condition, $F(1, 51)=6.41, p<.02$, as well as a significant main effect of choice dilemma, $F(1,51)=7.50, p<.01$ (see Figure 1). The main effect of choice dilemma reflects the fact that participants were more likely to choose the intuitive option for the apartment choice (51%) than for the ratio-bias paradigm (26%). The main effect of condition reflects the fact that, as predicted, participants in the fluent condition chose more intuitively (50%) than those in the disfluent condition (27%). The interaction was not significant, $F(1,51) = 1.02, p>.3$.




Figure 1a. Urn choices in Study 1.




Figure 1b. Apartment Choices in Study 1.

Figure 1. Participants' responses on the urn and apartment choices in Study 1.

Study 2: Speed of Thought

The results of Study 1 indicate that an internal state associated with a meta-cognitive sense of careful, deliberate processing leads people to favour rational analysis when faced with intuition/reason conflicts. According to the state cuing hypothesis, any internal state characterized by features strongly associated with intuition and reason—not just a sense of effortful processing—should similarly encourage individuals to trust one source of input over the other. For example, another widely-shared belief about intuitive impulses is that they arise rapidly and holistically, with little or no awareness of the processes by which they were derived (Inbar, Cone, & Gilovich, 2010; Nisbett & Wilson, 1977; Wilson, 2002). Reasoned assessments, in contrast, are seen as typically proceeding slowly and sequentially. Kahneman (2011), for example, summarizes this view by characterizing intuitive judgment as “thinking fast” and rational analysis as “thinking slow.”

According to the state cuing hypothesis, then, when people sense that they are thinking relatively rapidly, they should be relatively inclined to trust their intuitions; when they have the sense that they are thinking at a more measured pace, they should tend to trust their rational analyses. To test this idea, I induced participants to experience their thoughts as proceeding either quickly or slowly (Pronin & Wegner, 2006). I then presented them with two intuition/reason conflicts: (a) a version of the ratio-bias paradigm, and (b) a

hypothetical choice between two apartments. I predicted that participants who encountered these two choices after having been induced to think at a relatively rapid pace would be more inclined to trust their intuitions, whereas participants who were thinking at a slower pace would be more inclined to trust their rational judgment.

Method

Forty-one Cornell undergraduates participated in exchange for academic credit or \$5. Participants were seated in front of a computer screen and were told by the experimenter (who was unaware of the hypothesis) that they would be asked a series of questions, each with two possible answers. They were told that each question would appear on the screen and that they should read each one aloud before indicating their response verbally to the experimenter. They were further told that their answer to one of the questions would have real monetary consequences—that they could win \$20 at the end of the study depending on their choice and the outcome of a drawing.

Participants then read aloud and answered 12 trivia questions (e.g., “Which river is longer: the Hudson or the Potomac?”). These questions were presented using PowerPoint software, appearing on the screen in “typewriter” fashion (i.e., introduced one character at a time), and advancing at a pace designed to induce fast or slow processing (Chandler & Pronin, in press; Pronin, Jacobs, & Wegner, 2008; Pronin & Wegner, 2006). Through pre-testing, I determined that it took undergraduates reading at their own pace an average of

approximately 8 seconds to read each trivia question. To induce fast and slow thinking, I reduced this speed by 50% in one condition and increased it by 50% in the other. Thus, in the *fast thought* condition (n=20), the slides of the PowerPoint presentation automatically advanced every 4 seconds, whereas in the *slow thought* condition (n=21), the slides automatically advanced every 12 seconds.

After answering the 12 trivia questions, participants were then faced with the same ratio-bias and hypothetical apartment choices as those used in study 1. (The order of these two questions was counterbalanced.) These decisions were presented at the same speed for all participants (8 seconds per slide, again, typewriter fashion) to control the amount of time participants had to process the presented information. Although presented at the same speed, past research has shown that presenting the prior questions rapidly induces a sense of the mind working quickly and energetically on material immediately following, whereas presenting the prior questions slowly induces a sense of the mind working slowly and deliberately (Pronin & Wegner, 2006).

The PowerPoint slide on which the ratio-bias paradigm was presented read at the top in large font “\$\$\$ Real Money Decision \$\$\$” to indicate to participants that the drawing would actually take place and that they could win \$20 if a winning marble was drawn from their chosen urn. After the participant answered these two questions, the experimenter pulled out two urns – one containing 10 marbles, with 1 winner and 9 losers, and the other containing 100

marbles, with 9 winners 91 losers– and randomly drew a marble from the participants' chosen urn. If a winning marble was drawn, the participant was paid an additional \$20 for participating. If a losing marble was drawn, they received only their compensation for participating in the study (\$5 or academic credit). Participants were then debriefed, thanked, and dismissed.

Results

To assess the extent to which participants were inclined to follow the dictates of a rational analysis or side with a gut feeling, I performed a 2 (thought speed: fast or slow) × 2 (choice dilemma: apartment or ratio-bias) × 2 (choice order: urn choice first or apartment choice first) mixed-model analysis of variance with thought speed and order as between-subjects factors and choice dilemma as a within-subjects factor. This analysis yielded only the predicted significant main effect of condition, $F(1, 37)=5.83, p<.025$, which reflects the fact that participants who had to process the lead-in material quickly were more likely to choose the intuitive option (22%) than those who processed that material at a more leisurely pace (52%) (see Figure 2). None of the other main effects or any of the interactions were significant, all $F_s<1$. Thus, participants' fast or slow thought speed cued them to trust their intuitions or rational analyses, respectively.

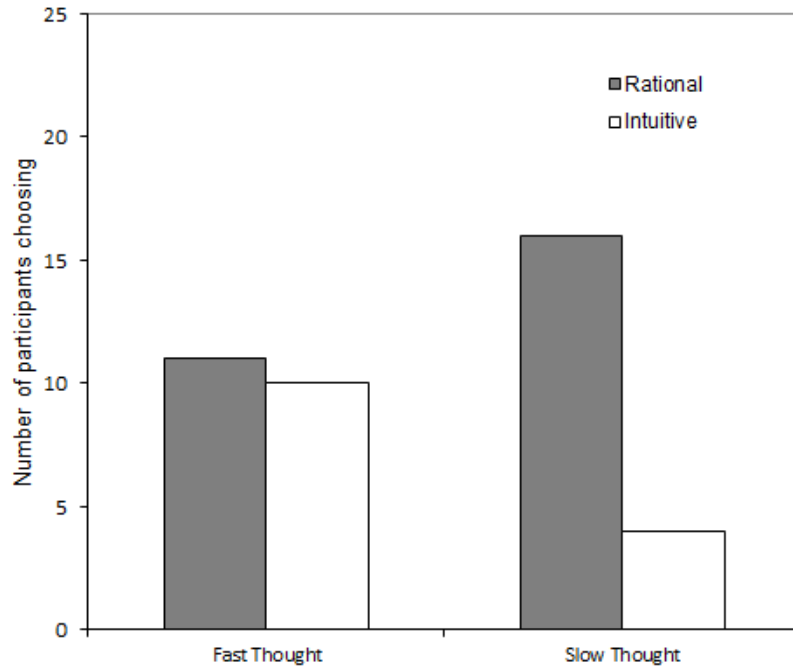


Figure 2a. Urn choices in Study 2.




Figure 2b. Apartment Choices in Study 2.

Figure 2. Participants' responses to the urn and apartment choices in Study 2.

Study 3: Inside/Outside

In Studies 1 and 2, I obtained consistent support for the state cuing hypothesis: Participants who were in mental states characterized by features typically associated with intuitive processing were more likely to opt for the intuitive option in an intuition-reason conflict, and participants in mental states characterized by features associated with rational processing were more likely to opt for the rational option. In Study 3, I examined the impact of a different association that people have to intuition and reason.

Specifically, I was interested in the feeling that intuitions seem to come from somewhere inside of us whereas rational judgment seems to exist outside of the self. The rules governing rational judgement, unlike intuitions, are learned from others, derived from explicit pre-existing principles, and often executed externally—as on a blackboard, a pad of paper, a calculator, or a computer. Because of these intuition-internal and reason-external associations, the state cuing hypothesis predicts that anything that makes a decision feel as though it is happening “in close” to the body should encourage intuitive responses and anything that makes a decision feels as though it is happening “at a distance” should encourage rational responses. I therefore manipulated the perceived distance of the decision by placing the two response alternatives on separate clipboards and having participants either hold the clipboards in their hands while they made a decision or observe the clipboards at a distance while they made their choice.

Method

Sixty-six Cornell undergraduates took part in the study in exchange for academic credit. Participants arrived individually and were seated at the end of a desk. To manipulate whether participants felt as though the decision was happening “inside” or “outside,” I put the decision-relevant materials on two clipboards. One of the clipboards had a piece of paper attached to it that read, in 72-size Calibri font, “Apartment A: Favored by Gut Feeling.” The other had a piece of paper attached to it that read, in the same font, “Apartment B: Favored by Rational Analysis.” Participants assigned to the *inside* condition were asked to hold the clipboards, one in each hand. Participants assigned to the *outside* decision saw the same clipboards as those assigned to the *inside* condition but, rather than considering the options while holding the clipboards in their hands, the clipboards were instead propped upright at other side of the desk, 4 feet away. (The left-right positioning of the two clipboards was counterbalanced.)

As soon as the participant was seated, the experimenter, who was unaware of the hypothesis of the study, either placed the clipboards in the participants’ hands (in the *inside* condition) or propped them up at the end of the desk (in the *outside* condition). He then told the participant: “We’d like you to imagine that you are looking for a new apartment and you have found two (A and B) that pique your interest—and that you are now having trouble deciding between them. Over the course of making your decision, you’ve noticed that on some attributes, Apartment A appears to be a better choice than Apartment B.

On others, Apartment B appears to be the superior choice. In the end, you are torn: your ‘gut’ is telling you to pick Apartment A, but when you reason everything out, comparing the two apartments on each attribute and weighing the relative importance of each attribute, your ‘head’ is telling you to pick Apartment B.” He gestured to each of the clipboards as he described the scenario. Participants indicated their choice to the experimenter orally. The experimenter then took the clipboards and debriefed, thanked, and dismissed the participant.

Results

Fifty-nine percent of participants in the *inside* condition chose the apartment said to be favoured by intuition, whereas only 31% of participants in the *outside* condition did so, $X^2=5.06, p<.05$.¹

General Discussion

Intuitive judgment has attracted a great deal of scholarly attention in recent years, resulting in many new and unexpected insights into the nature of snap judgments, the automatic components of judgment and choice, and why people are likely to second-guess what seems like the best choice from a rational

¹ There was also an unpredicted interaction between the order of the alternatives and the inside/outside manipulation, $X^2=12.00, p<.001$. Breaking down this interaction, the hypothesis was strongly supported when the intuitive choice was presented on the left: Fifty-seven percent of participants in the *inside* condition chose the apartment said to be favored by intuition, whereas only 7% of participants in the *outside* condition did so. However, when intuitive choices were presented on the right, this difference was substantially reduced: 62% chose the apartment favored by intuition in the *inside* condition, whereas 53% did so in the *outside* condition. Although highly speculative, this may have been driven by an odd feature of the situation in which intuitions were presented on the right, which was that participants were required to read from right to left, seeing choice “B” presented as the choice on the left, serving to make participants feel as though something was “amiss.”

perspective. Despite these recent insights into the nature and the power of intuitive judgment, however, we still know very little about how people figure out what to do when faced with a situation in which reason tells them one thing and intuition another. In this chapter, I sought to address this gap by examining one factor that may have an especially powerful influence on how people resolve such conflicts: the extent to which their internal mental state at the time of a decision resembles the features commonly associated with intuitive or rational thought—a phenomenon I have called *state cuing*.

Across three studies, I found that people seem to be cued by their current mental state to favour intuition or reason. When people are in mental states associated with deliberative processing—such as when they find it subjectively difficult to process information about a decision (Study 1)—they tend to trust the dictates of the pertinent rational analysis. Similarly, when people feel as though their mind is racing (Study 2) or that a decision is occurring “up close” rather than “at a distance” (Study 3), they are cued to “go with their gut” and to side with their intuitions.

The key feature in each case is that the mindset that people experience at the time of rendering a decision is one that closely matches the features of intuitive or rational thought. Because intuition is rapid and efficient, mindsets that induce rapid thought serve to encourage intuitive processing and thus make intuitive choices more compelling. Similarly, because reason is slow, deliberate, and effortful, a sense of subjective difficulty in processing information relevant

to a decision prompts a more rational orientation and makes rational choices more compelling. Across the three studies reported here, participants were induced to favour intuition or reason in both hypothetical decisions in which an intuition-reason conflict was made explicit (i.e., imagining a situation in which one is torn between two alternatives for an apartment rental; Studies 1, 2, and 3) and in real decisions with monetary consequences (i.e., the ratio-bias paradigm; Studies 1 and 2).

It is important to note that participants in my studies who favoured the reasoned analysis did not do so because they thought harder about the decision at hand and therefore were more likely to arrive at the normatively correct choice. In Studies 1 and 2, the rationally-optimal choice in the ratio-bias paradigm was obvious to all participants in all conditions because both urns were labeled with the chances of drawing a winner, facilitating an easy comparison between the two options. Furthermore, participants who were asked to choose between two hypothetical apartments did so in a context in which the intuition-reason conflict was made explicit and no deliberation was necessary to compute which option was favoured by reasoned analysis. I therefore believe that the most parsimonious explanation for the results across the studies is that the similarity between participants' internal states and the characteristics associated with reason led these participants to place more *trust* in the dictates of their rational analysis.

This raises a useful point of comparison with past theoretical accounts of how decision makers select choice strategies. These accounts have typically focused on two factors: the motivation to exert cognitive effort and the ability to do so (e.g., Chaiken, 1980; Payne, Bettman, & Johnson, 1993; Petty & Cacioppo, 1986). These factors surely play a large role in the selection of a decision strategy. However, the results reported here suggest that an exclusive focus on these two factors leaves out much of what affects people's choice of decision strategies. In the present studies, it is not the case that participants who made their choices deliberately did so because they were able and motivated to perform the effortful analyses required. Nor is it the case that those who chose intuitively did not know the rationally correct answer. Rather, participants were cued by their internal states to trust one of two already-available answers.

This phenomenon is consistent with previous research on task cuing in which we explored the predictors of how people resolve intuition/reason conflicts (Inbar et al., 2011) and it is also consistent with preliminary investigations of the state cuing hypothesis (Inbar, 2008). Just as this earlier research shows that people look *outward*—to properties of the decision task—to help them make a decision, the current work indicates that they similarly look *inward* to features of their current mindset or internal state at the time of the decision and use it as a diagnostic cue about which source of input to trust. In general, it seems that people are cued by available features of the surrounding

context—features both internal and external—that resemble those of intuitive or rational judgment and use these cues to help them to decide what to do.

These cuing processes may have implications for people’s satisfaction with their decisions and with their sense of themselves as decision makers. There are many occasions in which the outcome of a decision is not known until long after a choice is made. Was it right to choose Cornell over Princeton? Was it wise to pursue this field of study? Is she the one for me? In all of these cases, whether or not the decision was wise can only be known in the fullness of time. But does this mean that people put their evaluations of their decisions—or of themselves as decision makers—on hold? It seems unlikely. Some decisions “feel right” and others invite second-guessing long before the outcome is known. It may be that the fit between one’s internal state and the decision to choose the rational or intuitive exerts considerable influence on the sense that the choice was appropriate. So too might the fit between features of the decision context and the decision to choose rationally or intuitively (Inbar, Cone, & Gilovich, 2010). Because people tend to choose options that fit their internal states or the features of the decision context, they might tend to be rather satisfied with their choices, which might in turn feed the well-documented tendency for people to be overly confident in their judgments and decisions (Dunning, Griffin, Milojkovic, & Ross, 1990; Griffin, Dunning & Ross, 1990; Vallone, Griffin, Lin, & Ross, 1990). Conversely, those who choose in a way that runs counter to these cues might be especially prone to post-decision regret and dissonance—and so it

might be that going with a rational or intuitive option that “fits” leads to less dissonance, and less dissonance reduction, than choosing a rational or intuitive option that doesn’t match one’s internal state or the most prominent features of the decision context.

The current results may also have important implications for consumer choice. At auctions, for example—situations in which the auctioneer often rattles off information at a quick pace, thus inducing a sense of quick thinking and mania—individuals often report being “caught up in the moment” and losing their sense of appropriate rational considerations when making purchasing decisions, leading to greater impulse spending. Along the same lines, imagine online shoppers making their purchases either on a device on which decisions are made “up close” (say, an iPad or a mobile phone) or “at a distance” (say, on a desktop computer monitor or laptop that sits further away from the user). Many retailers (such as Amazon and Audible) offer mobile applications (e.g., for iPhones or Android-powered devices) that are equivalent to those that run on standard computers. Based on the results of Study 3, perhaps those who make their decisions using mobile applications are more inclined to trust their intuitions when deciding what to purchase than those who use versions of retail software that runs on devices for which the visual display is “at a distance.”

We often face difficult decisions in which the right path to take is unclear. In such cases, decision makers will often grab onto any useful input that they can to help them to decide what to do. Intuition-reason conflicts, by their

very nature, make for especially difficult choices, and here too, decision makers tend to grab onto any useful cues that are available, including what is happening inside their minds and inside their bodies.

**PART II: Prescription • When Can
Each Process Be Relied Upon?**

CHAPTER III. On the reliability of intuitive impulses

In Part I, I took an uncritical look at the way that people decide to trust the outputs of an implicit or explicit process in the context of intuitive judgment. However, in this descriptive analysis, I said very little about whether people are well-served by these meta-cognitive decisions to trust an intuition or not. In Part II, I now turn my attention to a prescriptive analysis of when implicit and explicit processes can serve as reliable sources for adaptive, satisfying decisions.

Of course, even though Part I was more focused on the predictors of people's choice strategies, some prescriptive inferences are perhaps unavoidable. After all, participants who ultimately came to trust their intuitions and chose the intuitively compelling urn in the ratio-bias paradigm did so at the expense of reducing their overall likelihood of leaving the experiment with additional compensation. In this way, the very nature of the paradigm is such that intuitive judgment is necessarily inferior to the normative standard by which the decision is judged. It is perhaps quite telling, then, that the cues I explored in the previous chapter still ultimately led individuals to favour their intuitions even in light of the rational standard.

I. When Are Intuitions Beneficial?

The ratio-bias paradigm serves as an excellent place to begin an analysis of when intuitions may or may not serve as a useful source of input into a decision. This decision dilemma is a quintessential example of what Hammond and colleagues (1987) referred to as an *indirect comparison*. In these types of

theoretical analyses, individuals' intuitive judgments about a problem are compared against a formal rational statistical or economic "rule" that serves as the standard for how the decision "should" be made. Hammond and colleagues have made the observation that such comparisons lead to the construction of situations that, by their very nature, can ultimately only serve to demonstrate that intuitive judgment is inferior to the rational standard—as evidenced by people's choices in the previous chapter—or equivalent to it, and never superior to it.

However, Hammond and colleagues have argued that such comparisons are not the appropriate metric by which to gauge the relative quality of intuitive judgment. Rather, to truly see the value of each choice strategy, researchers must focus instead on *direct comparisons*. In these types of comparisons, an individuals' intuitive judgment is compared against that same individuals' rational analyses of the problem, exactly as they occur in their actual decision-making.

This may seem like a subtle distinction, but it is important for at least two reasons. First, it emphasizes the disparities between reasoning "in the abstract" and reasoning as it actually occurs. Decision-makers vary quite substantially in their understanding of what constitutes an appropriate analysis or rational rule to apply to a given situation, and they often make mistakes in their application of these rules, particularly as rational analyses become increasingly complex and cognitively demanding. Second, the normative standard for a particular decision

often makes unreasonable assumptions about decision-makers' time constraints, motivation, and likelihood of possessing all of the information that is necessary to perform the analysis. Indeed, a number of researchers have proposed that the way that decision-makers ultimately decide—employing a number of “fast and frugal” heuristic shortcuts in their analyses—may be the optimal strategy when more reasonable assumptions about these kinds of constraints are made (e.g., Gigerenzer et al., 1999; see also Simon, 1955).

Through the lens of direct comparisons, researchers have, in recent years, begun to catalog a number of situations in which intuitions appear to gain an upper hand on rational analyses. Essentially all of these demonstrations have, either directly or indirectly, relied upon one particular strength of intuitive processes relative to more deliberative ones: their capacity. While intuitive processes can, by virtue of their parallel and holistic nature, usefully draw upon an essentially limitless amount of past experience, conscious deliberative processes are, by comparison, exceedingly limited in their capacity.

One prominent line of work that has honed in on this distinction is the naturalistic decision-making (NSM) tradition of research on intuitive judgment (e.g., Klein, 1999, 2003). Researchers in this tradition have proposed that one of the reasons why researchers focusing on heuristics and biases in judgment under uncertainty have found it relatively easy to construct experimental paradigms in which intuitions lead participants astray is that these demonstrations have focused exclusively on precisely the situations in which intuitions are likely to be

the least useful: that is, those in which intuitions cannot usefully benefit from past experience in a domain. On this view, to see the true power of intuitions, researchers must point their attention to domains in which individuals possess an extensive amount of past experience—what some researchers have called *domain-related knowledge* (Dane & Pratt, 2007). Indeed, when the intuitive judgments of experts in a particular domain are given greater consideration, intuitions appear to be substantially more beneficial, usefully taking advantage of environmental regularities (e.g., Brunswik, 1957; Simon & Chase, 1974; Hertwig, Hoffrage, & Martingnon, 1999; Klein, 1999, 2003).

Of course, if intuitions are to usefully benefit from past experience, then this past experience must be predictive and representative of future situations. Researchers in this tradition have thus sought to identify the quality of one's past experiences in a particular domain as a primary determinant of the relative value of an intuitive judgment. Hogarth (e.g., 2001, 2010), for example, has proposed that learning environments can be characterized as either *kind* or *wicked*. A *kind* environment is one in which feedback about the quality of one's intuitive judgments is both veridical and timely, thus allowing intuitions to become better-calibrated over time. A *wicked* environment, by contrast, is one in which feedback is scarce, misleading, unrepresentative, or otherwise biased. A quintessential example of a *wicked* environment, for example, is one in which one's experiences are based on an unrepresentative sample of possible future experiences, as when a doctor's patients with a particular set of symptoms

happen to always have a particular disease when, in fact, many other patients with the same set of symptoms happen to have other diseases.

Interestingly, other recent lines of work have suggested that intuitions can be relatively well-calibrated even in situations in which individuals receive no prior feedback about the quality of their gut-level reactions (e.g., Betsch et al., 2001; Dijksterhuis, 2004; Dijksterhuis et al., 2006). The paradigms employed by these researchers still seek to highlight the relative quality of intuitive judgment in situations in which individuals are exposed to an extensive amount of information that exceeds conscious capacity. However, these lines of work have posited that intuitive, associative processes can usefully capture some elements of past experience without intention or awareness and that these implicitly-learned elements of past experience can then be usefully applied to later decisions. For example, Betsch and colleagues (2001) have shown that when participants are incidentally exposed to an extensive amount of stock information running along a stock ticker on the bottom of a computer screen, individuals' gut-level reactions towards the stocks appear to be extremely well-calibrated to the stocks' actual performance—that is, their total value across all of the pieces of information encountered.

Thus, many recent lines of work converge on the notion that intuitions can usefully exploit prior experience and associative learning, thus giving intuitions an upper hand on rational analyses in situations in which: (a) intuitions have extensive (representative) prior experience on which to rely, and/or (b) the

information necessary to make a good decision exceeds the relatively limited capacity of conscious processes. However, there is a notable feature of these more recent lines of work. Although these paradigms do indeed overload participants with extensive amounts of information about decision alternatives, they nonetheless generally only expose participants to information that is relevant and necessary for an upcoming decision (e.g., Betsch et al., 2001; Dijksterhuis, 2004; Dijksterhuis et al., 2006).

This is perhaps an important artifact of these experimental designs because although intuitions may indeed usefully capture the breadth of one's past experience under these circumstances, there are often situations in which one's prior exposure is extraneous or invalid and must thus be discounted in order to make an adaptive decision. An important theoretical question thus concerns the extent to which the associative processes on which intuitive judgment is thought to rely possess a necessary filter for these sources of extraneous or invalid information.

Suppose, for example, that before participants were exposed to the extensive amount of stock information in Betsch and colleagues' (2001) paradigm, they happened to hear a rumor about one of the stocks that they later learn is fallacious. If the associative processes on which intuitions rely fail to properly discount such information, then intuitive judgments about the stocks will be systematically biased, even though all of the necessary, relevant

information needed for an accurate gut-level reaction was available to participants.

Indeed, some prominent theories concerning implicit phenomena have posited that one notable feature of associative processes is that they lack precisely this kind of validity filter (see, e.g., Gawronski & Bodenhausen, 2006, 2011), and can thus have important impacts on judgment and behaviour irrespective of the truth value of the information on which they are based. Thus, even if intuitions *can* usefully capture the breadth of one's past experience without feedback, the situations in which they can serve as a reliable source of information may be relatively more circumscribed than these recent demonstrations may suggest.

The goal of the current studies was thus to explore the extent to which all of one's past experience may be promiscuously captured by the associative processes on which intuitions are based, even in situations in which a subset of this information can be easily recognized (deliberatively) to be invalid or irrelevant—and, moreover, even when individuals have been exposed to all of the information that is necessary for an accurate, well-tuned intuitions.

That is to say, this was the original goal of the studies. However, to preview the findings, several attempts to replicate earlier research that had demonstrated that intuitions can implicitly capture elements of past experience without feedback came up short, making a proper test of the hypothesis impossible.

Study 4: Stock Ticker I

To properly examine the hypothesis that intuitions may sometimes capture irrelevant or invalid past experience, one must first establish that intuitive judgment is well-calibrated under learning conditions in which all of one's prior exposure is relevant, as in past research. The goal of Study 4 was thus to replicate the finding that intuitions can lead individuals to adaptive gut-level, spontaneous responses.

To this end, participants were first exposed to an extensive amount of information about the performance of a number of different stocks (see Betsch et al., 2001). While they saw this information and read it aloud, they were asked to pay attention to a series of print ads displayed on the top of the screen, which was meant to serve as a mild distraction during information acquisition.

Unbeknownst to participants, the performance of the stocks to which they were incidentally exposed varied markedly. After reciting all of the stock information, they were asked to provide their intuitive assessments of each stock, with the primary question of interest focusing on the calibration between participants' intuitive assessments and actual performance.

Method

Participants. Twelve Cornell undergraduates participated in exchange for \$3.

Stock ticker program. Participants were told that the study focused on their ability to memorize a series of rapidly changing print ads while they were

distracted by other information. To this end, they were told that while they were viewing colour images of print advertisements on the computer screen, a stock ticker would proceed along the bottom of the screen. Their task was to recite aloud all of the information that appeared on the stock ticker while simultaneously concentrating on the ads (see Figure 3).

The stock names that appeared on the ticker were similar to those used in past research (e.g., Betsch et al., 2001). On the basis of a short pre-test (n=70), I established that Cornell undergraduates had similar reactions to the stock names as previous (German) participants on two key dimensions: (a) perceived familiarity, and (b) intuitive judgments of the past performance of each of the stocks prior to any exposure.

The stock ticker information consisted of a total of 75 pieces of information—15 pieces for each of 5 stocks. Stock values were chosen to meet two criteria. First, the total value of each of the stocks—that is, the summation of all 15 pieces of information—was chosen so that they varied linearly in terms of their final performance (RONAT, 300; ELSKAR, 400; FAMO, 500; NARVIG, 600; PATEL, 700). Second, the minimum and maximum values for each stock were equated (+10 and +55, respectively) in order to prevent a single piece of extreme information from exerting an undue influence on people’s affective reactions.

The print ads consisted of a total of 60 high resolution colour images that cycled every five seconds on the screen. The stock ticker ran until the final stock value reached the left side of the screen, which took exactly five minutes. The

order of both stock information and print ads was randomized for each participant.



Figure 3. Screenshot of the stock ticker program. Colour print advertisements cycle every 5 seconds on the top portion of the screen while the stock ticker moves horizontally along the bottom portion.

Assessing intuitions. Following the stock ticker, participants completed an online questionnaire that assessed their feelings about each stock. To encourage participants to use their intuitions rather than rational analysis (see Betsch et al., 2001, Study 3), they were told that the next part of the experiment focused on the way that people’s intuitions operated and that, as a result, some participants would be asked to respond to the subsequent questions intuitively without thinking while others would be asked to think very carefully about their

answers. All participants were told that they had been assigned to the intuitive thinking condition, and should thus make their decisions spontaneously and as quickly as possible.

Next, participants saw each stock on a separate screen and were asked to report their feelings about it by clicking a location on a horizontal scroll bar that had anchors of -5 (very bad) and +5 (very good). Stocks were presented in a random order for each participant. Finally, participants were asked to provide demographic information and were then debriefed, thanked and dismissed.

Results

Table 1 summarizes participants' intuitive judgments about each of the five stocks as a function of actual performance. To discern whether participants' intuitions about each stock differed, I performed a one-way repeated-measured analysis of variance. This analysis revealed a significant main effect, $F(4,8)=8.314, p<.01$, suggesting that participants did indeed see distinctions among the stocks. However, in contrast to previous research (Betsch et al., 2001), these differences in participants' gut feelings failed to capture the actual performance of the stocks in any meaningful sense. By inspection, participants seemed to favour the lower-performing stocks relative to the higher-performing one, suggesting that they had very little insight into true performance. Indeed, in stark contrast to the perfect linear relationship between intuitions and actual performance observed by Betsch and colleagues (2001), the correlation between

participant’s intuitive judgments and the stocks’ actual performance was both quite modest and trending towards an inverse relationship, $r=-.217$.

Table 1. Intuitive judgments of stock performance as a function of actual performance.

	Stock Name (Actual Performance)				
	RONAT (300)	ELSKAR (400)	FAMO (500)	NARWIG (600)	PATEL (700)
Intuitive judgment (M)	.91	2.08	.50	.50	1.25
Intuitive judgment (SD)	1.83	1.62	2.47	2.32	2.26

Discussion

Thus, the current study failed to uncover any evidence that intuitive judgment could accurately summarize an abundance of numerical information encountered earlier in the experimental session. However, there were perhaps a number of features of the study design that may have unwittingly: (a) differed from previous research, and (b) had a substantive impact on participants’ intuitions in ways that served to reduce their accuracy. For example, one notable aspect of intuitive judgment that has been emphasized in a number of distinct lines of work is that it can be relatively easily contaminated by conscious deliberation, or “thinking too much” (e.g., Dijksterhuis, 2004; Dijksterhuis et al., 2006; Schooler & Engstler-Schooler, 1990; Wilson & Schooler, 1991). In this

study, it is perhaps not implausible that participants had some opportunity to consciously deliberate about the stock information in ways that could have contaminated an otherwise well-calibrated intuitive judgment. Because the stock ticker proceeded at a rather measured pace, with each piece of stock information appearing on the screen for a total of 15 seconds, there may have been an opportunity for subjects to recite each piece of information, direct their attention to the print ad, and still have some time to consciously mull over the stock information. Of course, subjects were told that their primary focus should be the ads, but this ultimately requires that participants carefully follow instructions—an assumption that may not be warranted (see Oppenheimer, Meyvis, & Davidenko, 2009).

Another possibility is that there may have been undue “noise” in the stock information acquisition process, disrupting, in some way, the associative processes that are thought to underly the capturing of this extensive prior experience. This could have occurred, for example, if any incidental sources of affect were unwittingly introduced while participants were encountering the stock information, thus swamping any effects of the stock values on participants’ intuitive impulses. Because many theoretical models of intuitive judgment have proposed that affect is a crucial feature of intuition (see, e.g., Damasio, 1994; Dane & Pratt, 2007; Epstein, 2008; Hogarth, 2001; LeDoux, 1996; Zajonc, 1980), we might expect that participants’ judgments could have been led astray by

these incidental sources of affect (Hogarth, Portell, & Cuxart, 2007; Hogarth, Portell, Cuxart, & Kolev, 2010; Weiss & Cropanzano, 1996).

A third possibility is that some theoretical models have posited that non-conscious processes require some time to operate on information acquired during past learning in order to consolidate it and better organize it in memory (Dijksterhuis, 2004; Dijksterhuis et al., 2006). Because participants in this study made their judgments immediately after finishing the stock ticker paradigm, there may not have been enough of a delay for such non-conscious processes to operate.

Finally, Betsch and colleagues (2001) have posited and provided evidence to suggest that participants' judgments will not be especially well-calibrated if subjects do not provide a spontaneous, intuitive judgment. One aspect of the original procedures that may have discouraged participants from doing so is that the precise numerical value of their responses on the dependent measures was prominently displayed to the left of the slider each time they clicked. Because some of our previous work on task cuing (Inbar, Cone & Gilovich, 2010; see also Chapter 2) has indicated that participants are relatively more inclined to trust the dictates of their rational analyses when a choice feels precise or numerically evaluable, this precise numerical feedback may have ultimately discouraged participants from trusting their spontaneous judgments, encouraging them instead to rationally deliberate about each of their judgments.

Honing in on these four possibilities, a number of changes were made to the design in the next study that sought to: (a) reduce participants' ability and motivation to consciously cogitate over the stock information, both during information acquisition and during the reporting of intuitive judgments, (b) reduce any potential extraneous influences that may have disrupted the acquisition of the stock information, and (c) introduce a delay between information acquisition and reporting of intuitions.

Study 5: Stock Ticker II

Method

Participants. Twenty-four Cornell undergraduates participated in exchange for \$3.

Procedure. Study 5 followed the same protocol as Study 4, except for the following changes. First, the speed of the stock ticker was increased by 20% as a means of reducing participants' ability to consciously deliberate on the stock information. This meant that participants were exposed to the 75 pieces of information about the stocks over four minutes rather than five.² Second, several of the advertisements that were judged to be potentially arousal- or affect-inducing were removed, as a means of reducing potential affective "noise" during the information acquisition process. This included, for example, an advertisement for a digital camera that featured a high resolution image of a

² It is also perhaps worth noting that this change necessitates another deviation from the procedures of the previous study. That is, participants saw only a subset of the total number of ads—48 rather than 60—and each participant saw a different subset of the ads, depending on the randomized order determined by the computer.

large, aggressively-poised snake; an advertisement that prominently featured a bald eagle and an American flag designed to evoke feelings of patriotism; and several advertisements that were designed to be humorous. These images were replaced with other, more affectively-neutral images so that there was still a total of 60 advertisements that cycled every 5 seconds. Third, a number of demographic (filler) questions were included at the beginning of the online questionnaire as a means of introducing a delay between the completion of the stock ticker and participants' reporting of their intuitions about each of the stocks. These included questions about participants' age, gender, major, and year in school. After answering these questions, they provided their intuitive judgments about each stock using the same procedures as Study 4, except that the response format was modified so that no numerical information was provided to participants while they made their judgments. Participants' responses were still recorded on a -5 to +5 scale as in the previous study. However, numerical feedback was never displayed to participants.

Results

Table 2 summarizes participants' intuitive judgments about each of the stocks as a function of actual performance. Submitting these responses to a one-way repeated-measures analysis of variance, participants in this version of the task failed to exhibit any significantly different gut-level reactions toward the stocks, $F < 1$, indicating that they had no strong insights into the stocks' true

values. Indeed, the correlation between intuitive judgments and actual performance was, once again, especially modest, $r=.36$.

Table 2. Intuitive judgments of stock performance as a function of actual performance.

	Stock Name (Actual Performance)				
	RONAT (300)	ELSKAR (400)	FAMO (500)	NARWIG (600)	PATEL (700)
Intuitive judgment (M)	0.52	1.19	0.20	1.16	1.04
Intuitive judgment (SD)	1.99	2.14	1.82	2.33	1.77

Discussion

It appears, then, that several theoretically meaningful changes to the study design did not serve to make intuitive judgments well-calibrated. However, these results are encouraging in the sense that participants in Study 5 would not be led astray by their intuitions as they would have been in Study 4. Whereas participants in study 4 had intuitively felt that the stocks with the lowest performance were subjectively the best, participants in Study 5 were merely indifferent to the stocks, exhibiting no strong inclinations towards any one of them. Perhaps, then, the changes made between these two designs were meaningful, but were not enough to prevent the various sources of bias that were resulting in miscalibration. The next study sought to test this hypothesis by

further reducing participants' ability to consciously deliberate over the stock information, and instituting an even large delay between information acquisition and reporting of intuitive judgments.

Study 6: Stock Ticker III

Method

Participants. Twenty-eight Cornell undergraduates participated in exchange for \$3.

Procedure. The procedures for Study 6 were identical to Study 5, except for the following changes. First, the speed of the stock ticker was further increased, such the total display time was 3 minutes, 20 seconds rather than 4 minutes. Second, to create a larger delay between the end of the stock ticker and the completion of the dependent measures, several additional questions were added to the filler questionnaire, including questions about participants' history in participating in psychology experiments, the location of their home town, and how often they had visited their family this semester.

Results

Table 3 summarizes participants' intuitions about each of the five stocks as a function of each stock's actual performance. As can be seen in the table, participants once again had no accurate insight into stock performance, nor any strong affective reactions to any of the stocks relative to the others, as confirmed by a one-way repeated-measures analysis of variance, $F < 1$. Indeed, the correlation between intuitive judgments and actual performance was, once

again, especially modest, $r=.22$. Thus, evidence for a well-calibrated intuitive judgment remained elusive under these experimental parameters.

Table 3. Intuitive judgments of stocks as a function of actual performance in Study 6.

	Stock Name (Actual Performance)				
	RONAT (300)	ELSKAR (400)	FAMO (500)	NARWIG (600)	PATEL (700)
Intuitive judgment (M)	1.01	0.94	0.84	1.08	0.87
Intuitive judgment (SD)	1.59	1.53	1.66	1.71	1.65

Discussion

It appears, then, that participants' intuitive judgments remained relatively miscalibrated in spite of attempts to reduce conscious deliberation and allow some time for non-conscious processes to operate on the information. There is, however, one additional possible explanation for the miscalibration of the previous three studies. Betsch and colleagues (2001) have suggested that one additional reason for instituting a short delay between the acquisition of information and the reporting of intuitions about that information is because of the possibility of a recency effect. On this view, intuitive judgments that are made immediately after acquiring the information are unduly influenced by the most recent pieces of information that were encountered about each decision

alternative and it is only with additional time that intuitions come to more closely reflect the breadth of one's past experience.

Of course, one reason to be skeptical of this possibility is that the length of the delay induced in each of the past experiments appeared to have relatively little influence on the quality of participants' intuitive judgments. If recency effects were indeed having an undue influence, it would be expected that longer delays should serve to improve the calibration between an intuitive judgment and actual performance. However, in the interest of thoroughness, I conducted one final replication attempt in which I more carefully controlled the last piece of information about each stock that participants encountered.

Study 7: Stock Ticker IV

Method

Participants. Fifteen Cornell undergraduates participated in exchange for \$3.

Method. The procedures for Study 7 were identical to Study 6, except that the five values that appeared at the end of the stock ticker were fixed (RONAT, +39; ELSKAR, +37; FAMO, +34; NARVIG +34; PATEL, +38). These values were chosen so that they represented those closest to the median of the distribution for each stock. The other 70 pieces of information were randomized as in the previous studies.

Results

Table 4 summarizes participants' intuitions about each stock as a function of its actual performance. As can be seen in the table, participants' intuitions were somewhat better-calibrated under this protocol than in the previous studies. However, they still nonetheless failed to show any strong intuitive pull towards any of the stocks, as evidenced by a non-significant one-way repeated measures analysis of variance, $F < 1$. The correlation between intuitive judgments and actual performance did show some evidence of improvement, $r = .69$. However, effects on par with the perfect correlation found in previous research—or, for that matter, marked differences in participants' intuitive judgments towards the stocks—remained elusive.

Table 4. Intuitive judgments of stocks as a function of actual performance in Study 7.

	Stock Name (Actual Performance)				
	RONAT (300)	ELSKAR (400)	FAMO (500)	NARWIG (600)	PATEL (700)
Intuitive judgment (M)	0.58	0.51	0.41	1.58	1.14
Intuitive judgment (SD)	2.07	1.85	1.54	1.82	0.98

General Discussion

In a number of lines of work, researchers have posited that intuitive judgment can implicitly capture many elements of one's prior experience with

decision objects, even when participants encounter an extensive amount of information and they receive no feedback about the quality of their gut-level reactions (Betsch et al., 2001; Dijksterhuis, 2004; Dijksterhuis et al., 2006). However, a notable feature of these lines of work is that participants are only exposed to the information that is necessary for an accurate, well-calibrated judgment. This raises the question of whether intuitions can remain relatively well-tuned even in the face of extraneous or invalid past experiences. In this way, although intuitions can indeed perhaps accurately capture the breadth of one's experience with a decision alternative, if the associative processes on which intuitions are thought to operate are overly promiscuous in the types of past experience that are captured, the situations in which intuitive judgment is beneficial and reliable may be relatively more circumscribed than current empirical demonstrations might suggest.

The current studies thus sought to examine the robustness of intuitive impulses to extraneous sources of information. However, across four attempts to replicate previous findings in which participants were exposed to objects that differed in their normative value, intuitions were found to be rather unhelpful even when individuals were only exposed to information that was relevant for an upcoming judgment. Indeed, in all four studies, participants exhibited very little insight into their previous experiences with stock information. This occurred despite several attempts to manipulate the extent to which participants were able to consciously deliberate over the alternatives—a factor that has been

claimed to influence the quality of intuitive insights (Betsch et al., 2001; Dijksterhuis, 2004; Dijksterhuis et al., 2006; Wilson & Schooler, 1991)—as well as eliminate extraneous noise in the information acquisition process. Because a well-calibrated intuition remained elusive, the hypothesis that intuitions can be relatively flighty in their development was ultimately untestable.

If we take seriously the dramatic differences in the results I have reported here and those of past work employing very similar procedures, there are at least three potential interpretations. First, intuitions may not adequately capture elements of past experience in the way that contemporary theories have claimed—at least the types of elements of past experience that are necessary for possessing a well-tuned intuition in the context of the stock ticker paradigm. This interpretation would imply that past results may have been the result of anomalies in the procedures or other statistical irregularities that gave rise to spurious conclusions. We cannot know for certain if this interpretation is correct, but it is a possibility that must be taken seriously in light of the fact that there are only a handful of studies employing this paradigm, most of which were conducted in a single laboratory (but see Plessner et al., 2008, for a conceptual replication of this result using markedly different procedures).

However, a more interesting theoretical possibility is that intuitions are indeed well-calibrated in situations like the one that the stock ticker paradigm seeks to capture, but there were important differences between the procedures I employed in these studies and those of past research—factors that served to

make intuitions inaccurate in my studies and accurate in previous ones. This is an interesting possibility because, if true, it has important implications for the assertions that motivated these replication attempts—that is, although intuitions are perhaps accurate under a certain, well-defined set of circumstances, this set of circumstances is considerably more circumscribed than one might expect.

From this perspective, there are at least two theoretically meaningful mechanisms that may have made intuitions quite accurate in one case and quite inaccurate in the other. First, unknown differences in the procedures may have given participants in my studies a greater opportunity to engage in deliberative thought, which contaminated an otherwise well-calibrated intuitive judgment. Betsch and colleagues (2001), for example, have found that if participants are forced to wait as little as six seconds before making their judgments, their intuitions become considerably less well-tuned to actual stock performance. More generally, as discussed earlier, a number of lines of work have sought to place the blame of poorly calibrated intuitions squarely on the shoulders of conscious processes that disrupt the weightings at which one arrives non-consciously (Dijksterhuis, 2004; Dijksterhuis et al., 2006; Wilson & Schooler, 1991). These studies highlight the notion that intuitive judgment is only valuable to the extent that it is unfettered by conscious deliberation. Of course, doubt is cast on this interpretation by the fact that manipulations that one would expect should have had an influence on the extent of participants' deliberation—that is,

the speed of the stock ticker—had very little influence on intuitive judgments across the replication attempts.

The other theoretically meaningful mechanism by which intuitive judgment could have been made less accurate is that procedural differences between these studies and past research had a direct influence on the quality of people’s intuitions—that is, some element of my implementation of stock ticker paradigm that prevented participants from properly acquiring or integrating the available information into an accurate summary evaluation. The implication of this possibility is that while intuitions may produce normative, accurate judgments, they can be affected quite substantially by some seemingly trivial, subtle factors that differentiate the procedures that I employed from those of past research. This is interesting in light of the fact that in both my procedures and those of past research, participants had complete access to all of the information that was necessary to make an accurate judgment. Yet it appears that they had relatively little access to this information when it came time to provide a gut reaction in my studies, and relatively good access to it in previous work. In any case, what seems clear from these replication attempts is that there is still a great deal of work needed to understand when intuitive judgments are accurate and well-tuned and when they are instead miscalibrated and biased.

CHAPTER IV. Changing your implicit mind: On the rapid formation and revision of implicit attitudes

In the previous chapters, I focused my attention largely on the nature of intuitive judgment. However, the notion that implicit processes can sometimes be a source of bias and error in guiding behaviour is an idea that pervades not just thinking about the relative value of intuition and reason but also theoretical models of many other phenomena. For example, research on stereotyping and prejudice has, in recent decades, emphasized the notion that individuals can harbour knowledge and evaluative responses toward outgroup members that are highly overlearned and thus acquire features of automaticity (e.g., Fazio, 2007). These rapid, spontaneous responses can sometimes be strongly inconsistent with one's deliberate assessments, yet because they are activated and sometimes applied relatively automatically, they can nonetheless exert a subtle but important, unintentional influence on behaviour. This is thought to be especially true in situations in which individuals fail to possess the motivation or ability to override these implicit responses, thus leading to behaviours that are sometimes inconsistent with more deliberate, conscious assessments (see Fazio, 2007; Hofmann, Gschwendner, Nosek & Schmidtt, 2005).

This is not to say, of course, that researchers have suggested that implicit responses are always maladaptive. Indeed, theorists have long proposed that implicit processes are highly adaptive in that they serve an important offloading function, guiding behaviours without conscious intervention and thus freeing up conscious resources for other tasks (see, e.g., Wigley, 2007). Thus, not unlike the

heuristics and biases program of research in judgment under uncertainty, implicit processes are generally thought to be adaptive and necessary, but they can nonetheless be an important source of unintentional and sometimes inescapable bias. Indeed, many theories emphasize that a primary role of conscious, deliberative processes is to serve as a monitor of the product of implicit processes, overriding these quick-but-sometimes-flawed responses when they may lead to error (Fazio, 2007; Gawronski & Bodenhausen, 2006, 2011; Kahneman, 2011; Wilson, Lindsey & Schooler, 2000).

Delving into theoretical conceptions of implicit and explicit processes in more detail, it becomes evident that the notion that implicit processes can be a source of error is predicated on a number of assumptions that have been made about their nature and operation. For example, theories of attitudes have drawn a sharp distinction between implicit and explicit attitudes, suggesting that they are governed by different processes and learning characteristics. Because implicit attitudes are generally thought to be governed by slow-learning associative mechanisms, their ability to incorporate newly-learned information is thought to be rather limited. Thus, one common and pervasive assumption about implicit attitudes that a number of prominent attitudes models share is that they are thought to be relatively insensitive to recent (but nonetheless relevant) learning and experience.

It is not difficult to see how such a slow-learning mechanism could ultimately be a source of bias and error in many situations. Although a number

of researchers have emphasized the value in an organism possessing a learning system or process that captures one's long-term experiences with an attitude object over time (e.g., Smith & DeCoster, 2000), there are many situations in which recent, relevant information should (normatively) be an important determinant of one's behaviour (Wilson, Lindsey & Schooler, 2000).

Suppose, for example, that one has just learned that one's romantic partner has been having an affair with one's best friend. This recently-learned information should arguably play a strong role in one's future attitudes and behaviours. Yet, because implicit responses are thought to be relatively insensitive to such information, they can thus sometimes encourage behaviours that are inconsistent with more deliberative assessments (that is, those that more easily incorporate recent learning experiences)—as anyone who has ever “drunk-dialed” their ex-boyfriend or -girlfriend knows quite well. A similar situation occurs when someone whom one likes and admires succumbs to a dramatic moral failing that tarnishes one's view of him or her. In each of these situations, this recent information should, perhaps, play a strong role if one is to engage in adaptive behaviour. The implication, then, is that implicit responses, by virtue of remaining relatively insensitive to these kinds of recently-learned yet highly relevant experiences, may need to be overridden in favour of a more explicit, deliberate response in precisely these types of situations.

Yet, despite the pervasiveness of this assumption about the nature of implicit attitudes, there is ultimately very little research that has attempted to

examine it empirically. In this chapter, my goal is thus to explore this assumption in more detail, asking whether implicit attitudes are indeed insensitive to new, relevant learning, and whether they must necessarily be a source of bias and error in people's implicit evaluations of targets in situations in which recent experience is relevant and important.

I. Gradual Development, Slow Change

Taking a more detailed look at theoretical models of implicit attitudes, the traditional assumption that they are slow to develop and resistant to change once established (e.g., Bassili & Brown, 2005; Conrey & Smith, 2007; Rydell & McConnell, 2006; Smith & DeCoster, 2000; Strack & Deutsch, 2004; Wilson, Lindsey, & Schooler, 2000; cf. Gawronski & Bodenhausen, 2006, 2011) is one that follows primarily from the commonly held perspective that their operation depends largely on associative processes. Because these processes are assumed to operate largely via the incremental accrual of information that slowly strengthens associations in memory over time (e.g., McClelland & Rumelhart, 1985; Smith & DeCoster, 2000), implicit attitudes are similarly thought to require extensive and consistent experience with stimuli to become established. Moreover, once an implicit response towards an attitude object has developed, a similar line of reasoning suggests that it is only through extensive exposure to countervailing information that this response can be revised or “undone”—a prediction that has gained some recent empirical support (Rydell & McConnell, 2006; Rydell et al., 2007). Indeed, this assumption explains why most published

attempts to change people's implicit attitudes have largely relied on evaluative conditioning procedures that provide exactly this sort of consistent and extensive countervailing information over many dozens or hundreds of trials of learning (e.g., Olson & Fazio, 2001; Rydell & McConnell, 2006).

But are these assumptions about the learning characteristics of implicit attitudes correct? In the attitudes literature, there has been surprisingly little research investigating the developmental trajectories of newly formed implicit attitudes, and the current empirical landscape is decidedly mixed. On the one hand, some work challenges contemporary theoretical assumptions by showing that implicit attitudes can form without extensive or repetitive experiences (for a review, see Gawronski & Bodenhausen, 2006). Several studies have shown that people appear to be able to implicitly evaluate novel stimuli on the basis of very little information acquired only moments earlier (Ashburn-Nardo, Voils & Monteith, 2001; Castelli, Zogmaister, Smith, & Arcuri, 2004; Gregg, Seibt, & Banaji, 2006; Otten & Wentura, 1999). For example, Ashburn-Nardo et al. (2001) found that participants exhibited an implicit preference toward an ingroup to which they have been assigned in a minimal group paradigm only a few minutes earlier. Similarly, Gregg et al. (2006) found that when participants were asked to suppose that novel, fictional groups possessed good or bad traits, the participants immediately exhibited implicit positivity toward the good versus bad group.

On the other hand, however, other research suggests that once an

implicit response has formed towards an attitude object, it becomes largely insensitive to revision thereafter, even if it was formed only moments earlier (Gawronski & Strack, 2004; Gregg et al., 2006; Rydell & McConnell, 2006; Rydell et al., 2007; Wilson et al., 2000)—a result that is more aligned with traditional as well as recent theoretical conceptions (see Gawronski & Bodenhausen, 2011, p. 88). In one of the few studies exploring the sensitivity of implicit attitudes under minimal conditions of re-learning (as opposed to extensive evaluative conditioning, as described above), Gregg et al. (2006) tested whether participants would be able to incorporate new, countervailing information about fictional groups into their recently-formed implicit attitudes toward those groups. They found that implicit responses were largely insensitive to exposure to new information, appearing to be “easier done than undone.”

Thus, empirical findings on the developmental trajectory of implicit attitudes are both scarce and decidedly mixed. What can explain such an unclear empirical landscape, in light of such strong theoretical claims? Most attitude theories characterize associative processing as (largely or only) involving semantic memory, which does appear to afford only slow, gradual associative learning. However, there are other kinds of associative processing (Amodio & Ratner, 2010; Poldrack & Foerde, 2007; Squire & Zola, 1996) which can occur extremely rapidly, even within a single trial. For example, instrumental, goal-directed learning as well as Pavlovian conditioning in animal models develops rapidly (e.g., Hermer-Vazquez et al., 2005; Yin & Knowlton, 2006), including

during single trial episodes (e.g., see Cahill & McGaugh, 1990; Hilliard, Nguyen, & Domjan, 1997). Similarly, animals and humans have been repeatedly shown to learn avoidance responses to an aversive stimulus after only a single encounter (e.g., LeDoux, 2000; Rutishauser, Mamelak, & Schuman, 2006). Although there is relatively little evidence establishing the kinds of associative processing underlying implicit attitudes, some researchers have argued against the sole involvement of semantic memory (see Amodio & Devine, 2006).

The view from cognitive psychology and neuroscience that multiple kinds of associative processes exist, coupled with the empirical evidence in many animal models for rapid learning, suggest that the notion that implicit attitudes may only be acquired through slow, gradual learning should be re-visited. In light of this work, I would expect that there are situations in which implicit attitudes can develop quickly (as recent work has found), and may also change rather quickly in light of new information (which has not yet been shown).

Why, then, has previous work failed to find evidence of such sensitivity to new information (e.g., Gregg et al., 2006; Rydell et al., 2006; Rydell et al., 2007)? It has been widely assumed for many decades throughout the memory and learning literatures that organisms are more responsive to learning about objects that are instrumental to their outcomes (e.g., “unconditioned stimuli” need to have baseline relevance for the learner). Thus, previous work on this question, by virtue of using hypothetical groups which have no immediate or future consequences for participants, may not have been a sufficiently sensitive test of

rapid revision. That is, although it may be relatively easy to form an attitude toward a novel object in an “evaluative vacuum,” when there is no strong pre-existing information, the revision of one’s attitude in the face of new information may require that the stimulus have some minimal relevance to the person.

In 6 experiments, I introduced participants to novel attitude objects that have been shown repeatedly in the social psychological literature to have instrumental value to most people, including members of one’s ingroup (in Studies 8, 9a, and 9b); a character in a video game that either helped or hindered participants’ performance (Studies 10 and 11); and a target individual who participants learned had been convicted of a serious crime (Study 12). After the introduction of the novel attitude object, later in each experiment, I provided new, countervailing information, assessing implicit attitudes both before and after this information was provided. My main objectives were to: (a) provide additional evidence for the rapid formation of implicit attitudes, and, more importantly, (b) provide an existence proof that people are able to revise newly-formed implicit attitudes on the basis of minimal re-learning.

Study 8: Minimal Group I

Method

Overview. Participants were randomly assigned to a novel ingroup using a minimal group paradigm and their implicit attitudes towards members of the ingroup and outgroup were assessed. They were then led to believe that there had been a mistake in their feedback and that they actually belonged to the

other group. Participants were asked to re-do the implicit attitudes measure while keeping their true assignment in mind, with changes in implicit responses from time 1 to time 2 assessed.

Participants. Eighty-three undergraduates participated in exchange for academic credit. Three participants were excluded due to computer malfunctions; four were excluded because they failed to realize that the fake data file (described below) contained inconsistent information; and three were excluded because they made a large number of errors on the second IAT (>10%). This left a final sample of 73 participants.

Procedure

Introducing the novel stimulus. To assign participants to a novel ingroup, I used a procedure developed by Ashburn-Nardo, Voils, & Monteith (2001). Participants were told that the computer would assess their artistic preferences by having them rate how much they liked a collection of 20 paintings painted by two different artists—*Quan* and *Xanthie* (see Ashburn-Nardo et al., 2001, Study 3). Afterwards, they received (randomly-assigned) false feedback that indicated that they had a strong preference for paintings created by the artist *Xanthie* (*Quan*) by a margin of 90.3% to 9.7%.

Measuring implicit attitudes. After ingroup assignment, participants' implicit attitudes towards members of each group were assessed using an Implicit Association Test (IAT; see Greenwald, McGhee, & Schwartz, 1998). Participants were told that they would be categorizing a list of peoples' names

on the basis of which of the two artists each person preferred. The names were unfamiliar to participants³, but they were told that a simple heuristic could be used to determine which artist the target person preferred: those names containing a *q* should be categorized as preferring *Quan*, whereas those containing an *x* should be categorized as preferring *Xanthie*.

The IAT consisted of 7 blocks of 40 trials each. Only data from the critical blocks (that is, blocks 3, 4, 6 and 7) were analyzed. In 2 of the critical blocks, participants categorized the targets as either *Quan or Unpleasant* or *Xanthie or Pleasant*. In the other 2 critical blocks, participants categorized targets as either *Quan or Pleasant* or *Xanthie or Unpleasant*. (The order of these critical blocks was counterbalanced. Order did not interact with any of the independent variables and is thus not discussed further.)

Reversing the meaning of the novel stimulus. After completing the IAT, the experimenter entered the cubicle and indicated that she “needed to check something.” What happened next differed by condition. In the *switch* condition, the file contained what looked like the participant’s data from the painting preferences task, including two lines at the bottom of the file that were the purported summation of their preferences across the 20 paintings (see Figure 4). However, this file was modified for each participant to have the opposite feedback from what participants originally received. When the participant

³ For this task, I used the same stimuli employed by Ashburn-Nardo, Voils, & Monteith (2001): **Q names:** Qutar, Aloqu, Boqer, Moraq, Narqo, Quy, Qazi, Qesh, Linquos, Jaqer. **X names:** Bixten, Volx, Nardoxa, Xerdo, Xercerla, Maxson, Yexx, Xamieh, Xindau, Lixner.

noticed the error, the experimenter apologized, explained that the computer had done this a few times today, and asked if they would be willing to re-do the task so that their data could be salvaged.

OVERALL TOTAL FOR QUAN:	9200	90.7%
OVERALL TOTAL FOR XANTHIE:	1280	9.3%

Figure 4. The last two lines of participants' purported data file for a participant initially assigned to prefer *Xanthie*.

In the *control* condition, the file contained a series of headers where it appeared as though data should have been recorded. However, the rest of the file was blank. After noticing the data recording error, the experimenter apologized, explained that the computer had done this a few times today, and asked if participants would be willing to re-do the task so that their data could be salvaged. This condition was included to assess the effects of completing the IAT a second time without a switch in group assignment.

After participants completed the IAT a second time, they were debriefed and thanked. (No participants reported any suspicion in either condition.)

Results

To create a measure of implicit attitudes towards *Quan* and *Xanthie*, I calculated two D-scores for each participant (see Greenwald et al., 2003). Each D-score represented the direction and magnitude of participants' implicit

preference for one target group over the other either at time 1 or time 2. The score is calculated based on the differences in participants' reaction times in the critical blocks of the IAT. In this case, positive D-scores reflect an implicit preference for *Quan* over *Xanthie*, whereas negative D-scores reflect the reverse.

To assess the effects of the manipulations on participants' implicit attitudes, I performed a 2 (initial ingroup assignment: *Quan* or *Xanthie*) × 2 (switch condition: switch or control) × 2 (time: 1 or 2) mixed-model analysis of variance in which group assignment and switch condition were between-participants factors, and time was a within-participants factor. This analysis yielded the predicted 3-way interaction, $F(1,65)=4.992$, $p<.05$ (see Figure 5).

To understand the nature of this interaction, I conducted follow-up simple effect analyses. First, to assess whether participants' quickly formed an implicit attitude towards people in their ingroup, I compared time 1 implicit attitudes for those initially assigned to *Quan* to those initially assigned to *Xanthie*. This analysis yielded a significant main effect of group assignment at time 1, such that those who had been assigned to *Quan* showed a greater preference for *Quan* than those assigned to *Xanthie*, $F(1,71)=18.195$, $p<.001$.

Next, to assess whether participants rapidly revised their implicit responses in light of the "mistake" in their feedback, I analyzed responses separately for those in the *switch* condition. This analysis yielded the predicted 2-way interaction between time and initial group assignment, $F(1,30)=11.509$, $p<.01$. Breaking down this interaction further, I examined changes in

participants' implicit preferences from time 1 to time 2, separately for those initially assigned to *Quan* and *Xanthie*. This analysis revealed that those who were initially assigned to *Quan* exhibited a significant shift towards preferring *Xanthie* at time 2, $t(19)=2.695, p<.05$. Similarly, participants in the *Xanthie* condition exhibited a significant shift towards preferring *Quan* at time 2, $t(18)=2.273, p<.05$. In contrast, in an analysis of the 2-way interaction for those in the *control* condition, the only significant effect was an (unpredicted) main effect of time, $F(1,35)=6.882, p<.05$, such that attitudes became more neutral at time 2.




Figure 3a. Implicit Responses in Study 8 in the control condition. Positive values reflect a preference for Quan relative to Xanthie whereas negative values reflect the reverse.




Figure 3b. Implicit Responses in Study 8 in the switch condition.

Figure 5. Implicit Responses (D-Scores) towards Quan/Xanthie in Study 8.

Discussion

When participants were led to believe that their group assignment had been based on mistaken feedback, their implicit responses towards each of the groups exhibited significant changes that were consistent with this newly-acquired information. However, when participants were led to believe that their data had merely failed to record properly, their implicit responses exhibited similar changes independent of their initially assigned group, merely becoming more neutral over time. These results provide preliminary evidence that participants' implicit responses towards novel targets can indeed develop and change relatively rapidly, even on the basis of minimal information (cf. Gawronski & Strack, 2004; Gregg et al., 2006; Wilson et al., 2000).

However, a prerequisite for such rapid revision may be that the novel stimulus has some relevance to people. Indeed, one notable difference between the procedures I employed here and those of past research is that participants in past work were asked to consider hypothetical groups that were of little relevance to them and had no consequences for their future outcomes or behaviour. In contrast, in the current studies, participants implicitly evaluated an *ingroup*—and changes in such a group assignment were thus of considerably greater consequence to them (Campbell & Tesser, 1985; Tesser, 1988).

One important caveat to this claim, however, is that participants in the control condition also exhibited an (unexpected) effect of time between the first and second assessment of their implicit attitudes. Because of the nature of the

construction of the IAT—that is, because implicit responses are measured by having participants’ categorize items as part of one group *or the other*—the interpretation of this effect of time is somewhat ambiguous. It could be that participants are exhibiting an inclination towards greater positivity towards *Xanthie*, an inclination towards greater negativity towards *Quan*, or some combination of these two possibilities.

It is also important to note that the interpretation of participants’ time 2 attitudes is similarly unclear. Although participants exhibited a significant shift in their preferences from time 1 to time 2, they nonetheless failed to show any strong preference for one group over the other at time 2. This result raises the possibility that participants’ time 2 attitudes may reflect “noise” or implicit ambivalence toward the groups (e.g., Petty, Tormala, Brinol, & Jarvis, 2006) rather than a clear shift in their preferences that cleanly reflects most recent learning. However, given that participants in three of the four cells exhibited significant shifts towards a preference for *Xanthie*, the fact that participants in the *switch* condition who were initially assigned to *Xanthie* showed a significant shift towards preferring *Quan* is that much more impressive.

Studies 9a and 9b: Minimal Group II

Studies 9a and 9b were conceptual replications of Study 8. However, in an attempt to address the shortcomings of Study 8, a number of changes were made to the procedures. First, the IAT was replaced with a different implicit measure—the Affect Misattribution Procedure (AMP; see Payne et al., 2005).

This procedural change was introduced with the goal of ensuring that the completion of two consecutive implicit measures was relatively short, easy, and less tedious. This change also necessitated a change in the way that ingroup and outgroup members were identified in the implicit measure (that is, the heuristic that participants used to distinguish members of each group)—these differences are outlined in more detail below.

Second, I used a different procedure to minimally assign participants to an ingroup, using a different arbitrary criterion for the assignment and different group names. This procedural change was used to ensure that the results of the previous study were not an idiosyncratic feature of the group names or minimal group assignment procedure employed in the previous study.

Method

Overview. As in Study 8, participants were randomly assigned to a novel ingroup using a minimal group paradigm and their implicit attitudes towards members of the ingroup and outgroup were assessed. They were then led to believe that there had been a mistake in their feedback and that they actually belonged to the other group. Participants were then asked to re-do the implicit attitudes measure while keeping their true assignment in mind.

Participants. Experiment 9a: Ninety undergraduates participated in exchange for academic credit or \$4. Four participants were excluded due to computer errors, and 2 were excluded because the experimenter accidentally

provided the wrong false feedback in the “mistake” paradigm (described in more detail below).

Experiment 9b: Fifty-nine undergraduates participated in exchange for academic credit or \$4. Two participants were excluded due to a lack of variance in their responses on the implicit measures (described in more detail below).

Procedure

Introducing the novel stimulus. To assign participants to a novel ingroup, I used a procedure developed by DeSteno, Dasgupta, Bartlett, and Cadjric (2004). Participants were asked to make a series of estimations (e.g., “How many people ride the New York subway each day?”), and given (false) feedback about their likelihood of overestimating or underestimating quantities, resulting in their assignment to one of two groups. In Study 9a, I used the same group names used by DeSteno et al.: *Underestimators* or *Overestimators*. In Study 9b, I used the novel nonsense words *Radmore* and *Markens*. For Study 9b, on the basis of two short pre-tests, I established that the group names *Radmore* and *Markens* were equivalent in terms of participants’: (a) explicit familiarity and liking of the group names, and (b) implicit evaluations of the group names prior to any exposure to them. To help participants remember their group assignment, groups were assigned the colour green or blue (counterbalanced), and participants were asked to confirm the colour of their group on a short questionnaire.

Measuring implicit attitudes. After ingroup assignment, participants' implicit attitudes towards people belonging to each group were assessed using an Affect Misattribution Procedure (AMP; see Payne et al., 2005).

Before the task began, participants were shown pictures of 6 ingroup members and 6 outgroup members. The image of each face was tinged to reflect their group assignment (6 green, 6 blue) and the text of the group assignment was indicated below each face (e.g., "Underestimator"; "Radmore")

Next, participants completed the AMP. Each trial consisted of: (a) a prime (75 ms), (b) a Chinese pictograph (100 ms), and, (c) a backward mask (random noise), displayed until participants responded. Participants' task was to indicate whether they thought the pictograph was more or less pleasant than average. Payne and colleagues (2005) have shown that people misattribute their automatic evaluation of the prime to the pictograph, providing an unobtrusive measure of people's implicit attitudes toward the prime.

I used a total of 60 primes: 20 blue-tinged faces, 20 green-tinged faces, and 20 neutral images (grey-scale squares, pictures of office supplies, etc).

Participants responded by pressing *d* or *k* on the keyboard.

Reversing the meaning of the novel stimulus. After completing the AMP, the experimenter opened a file containing what looked like the participant's data, including 4 lines that were the ostensible summation of their responses (see Figure 6). However, this file was prepared for each participant so that it contained the opposite feedback from what they originally received. Once the

participant noticed the error, the experimenter apologized, explained that the computer had done this a few times today, and asked if they would be willing to re-do the task so that their data could be salvaged (see Gregg et al., 2006). After they agreed (all did so), they indicated their new group colour on a questionnaire, and completed the AMP a second time. They were then debriefed and thanked. (No participants reported any suspicion of the manipulation.)

OVERALL TOTAL FOR ACTUAL:	7880
OVERALL TOTAL FOR ESTIMATE:	9220

DIFF SCORE BY WEIGHTING:	-1340
ASSIGNMENT:	OVERESTIMATOR

Figure 6. The last four lines of participants' purported data file, in this case for those initially given feedback that they were an Underestimator (in Study 9a).

Results

To create a measure of participants' implicit preference for each of the group targets, I calculated the proportion of times participants indicated that the target pictograph was more pleasant than average separately for each target group (*Underestimators* and *Overestimators* in Study 9a; *Radmore* and *Markens* in Study 9b). Next, I calculated a difference score between the proportions of

times participants chose *more than pleasant* for each of the groups. This measure thus reflects an implicit preference for one of the groups relative to the other.⁴

Because the studies were identical except for the names of the novel groups, I combined the data sets to increase power, and included the variable of Experiment in my analyses.⁵ I then performed a 2 (initial assignment: Novel Group 1 or 2) × 2 (ingroup color: green or blue) × 2 (time: 1 or 2) × 2 (Exp. 9a or 9b) mixed-model analysis of variance. (Note that, in this analysis, Experiment was not randomly assigned.) This analysis yielded the predicted 2-way interaction between time and initial assignment, $F(1,129)=8.394$, $p<.01$ (see Figure 7), and this interaction was not qualified by any effects of Experiment (all $ps > .49$), indicating that the differences in the familiarity of the group names did not influence the pattern of results.⁶

To understand the nature of this interaction, I conducted simple effects analyses. First, to assess whether participants quickly formed an implicit preference for their ingroup, I compared participants' responses at time 1 for

⁴ For Experiment 9a, positive values reflect an implicit preference for *Overestimators* relative to *Underestimators*. For Experiment 9b, positive values reflect an implicit preference for *Radmore* relative to *Markens*.

⁵ Note that if Experiment moderated any effects, it would not be possible to determine whether this was driven by the different novelty of the group names because of a lack of random assignment. However, these changes in the experimental design appear not to have had any substantive effect on the dependent measures.

⁶ To ensure that the effects in these studies were not the product of the way that the data from the two experiments were combined, I conducted the analysis in two ways. In the first analysis, I assigned Novel Group 1 to *Overestimator* (in Study 9a) and *Radmore* (in Study 9b). In the second analysis, I assigned Novel Group 1 to *Overestimator* (in Study 9a) and *Markens* (in Study 9b). The results were equivalent between these two analyses, indicating that the results are not the product of the particular way that groups were assigned to Novel Group 1 or Novel Group 2.

those initially assigned to Novel Group 1 to those initially assigned to Novel Group 2. This analysis revealed a stronger implicit preference for Novel Group 1 when it represented participants' ingroup than when it represented participants' outgroup, $F(1,135)=10.713$, $p<.001$. Next, to assess whether participants rapidly revised their implicit preferences in light of the "mistake" in their feedback, analyzed the effect of time separately depending on participants' initial group assignment. This analysis revealed significant shifts in participants' implicit preferences when they were initially assigned to Novel Group 1, $F(1,135)=2.913$, $p<.05$ (one-tailed), as well as when they initially assigned to Group 2, $F(1,135)=6.829$, $p<.01$ (one-tailed). However, there were no significant differences of participants' responses at time 2, failing to uncover evidence of a full reversal, $F<1$.

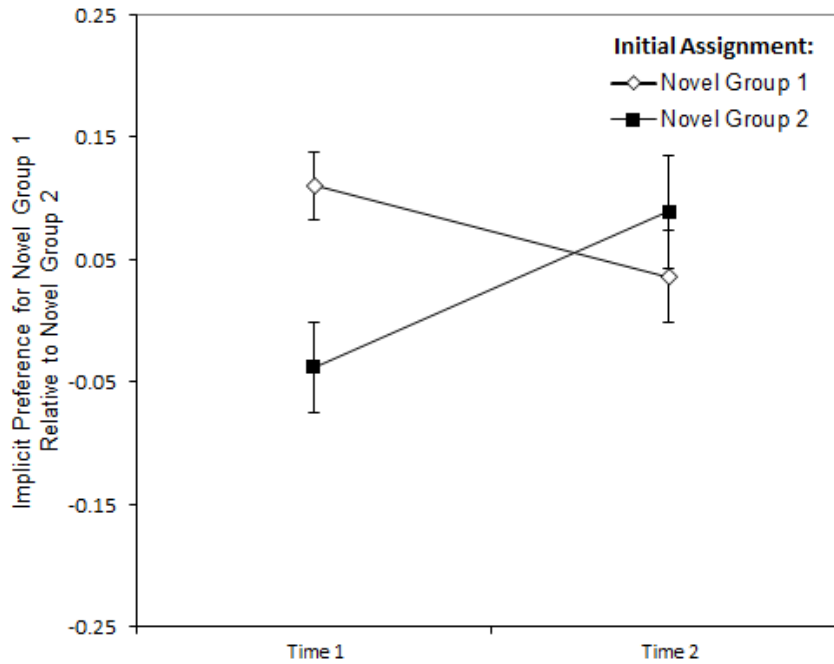


Figure 7. Participants' implicit preferences for novel group 1 over novel group 2 before (time 1) and after (time 2) learning of the "mistake" in their computer feedback.

Discussion

Conceptually replicating Study 8, participants once again exhibited significant shifts in their implicit preferences after learning of a "mistake" in the feedback they received about their group assignment. This result provides additional evidence that is consistent with the assertion that implicit attitudes can change quite rapidly under certain conditions. Like the previous study, one of the potential explanations for this rapid revision, I propose, is that the novel stimulus employed in the current study was of relevance to participants by virtue of its connection to the self (Campbell & Tesser, 1985; Tesser, 1988).

Although these results are entirely consistent with rapid revision, because attitudes theories make strong predictions about the developmental trajectories of implicit attitudes—not to mention some evidence already in the literature that is inconsistent with this result (Gawronski & Strack, 2004; Gregg et al., 2006; Wilson et al., 2000)—I sought, in the next study, to conceptually replicate the results of the previous two studies in an entirely different domain that would nonetheless still be relevant to participants: a video game in which they interacted with objects that either helped or hindered their goal of achieving a high score.

Study 10: *Wugs I*

Overview

Participants were told that they would be playing two rounds of a video game in which they would be interacting with a novel stimulus. The meaning and relevance of the novel object shifted from one round of the game to the other. Whereas it was beneficial to the participant in one round (that is, afforded them points), it was harmful in the other round (that is, could steal all of their points and reset their score to zero). Participants played one round of the game (counterbalanced), and then completed a measure of their implicit attitudes towards the object. Next, they played the second round of the game and then completed a second measure of implicit attitudes.

Method

Participants. Ninety-nine Cornell undergraduates participated in exchange for academic credit.

Procedure

Introducing the novel stimulus and changing its evaluative meaning.

Participants played a short video game modeled after the classic game Pac-Man. In my modified version, players interacted with a novel object that replaced the traditional “ghosts” from Pac-Man. I labeled this object a *wug* (Berko, 1958), depicted as a half-white-half-blue triangle with eyes (see Figure 8).




Figure 8. A wug.

To examine any potential effects of the length of exposure to the attitude object, the amount of time that participants played the game was manipulated. In one condition, they played each round for thirty seconds before being stopped by the experimenter. In the other condition, they never had an opportunity to play the game. Instead, they merely read the instructions for how each round of the game would be played, with their implicit attitudes towards *wugs* assessed immediately after they read the instructions for each round. This condition

allowed for an assessment of how participants' implicit preferences responded to the expectation of an interaction with an attitude object cast in a positive or negative light.

After playing each round of the game (or reading the instructions for how the round would be played), I assessed participants' implicit attitudes towards *wugs*, assessing changes from time 1 to time 2.

Positive Round. In one round of the game, *wugs* were described as helpful. Participants were told that if they could touch the *wug* with their character (a small purple square that replaced Pac-Man from the original game) they would earn points.

Negative Round. In the other round of the game, the *wugs* were described as harmful and participants were told to avoid them at all costs. In this round, if the *wugs* touched the player's character, they would be killed and would lose a life.

Assessing implicit attitudes. Implicit attitudes were once again assessed using an AMP (Payne et al., 2005). The specifications for this task were precisely the same as in the previous study, except that the primes consisted of: (a) 20 target images consisting of images of *wugs* as well as the word "*wug*", and (b) 20 neutral images, consisting of various coloured squares, pictures of office supplies, etc., as well as the novel word "*dax*." Participants indicated their response by pressing either *d* or *k* on the keyboard.

Results

To assess whether participants' implicit preferences were sensitive to their most recent experience with the attitude object, I performed a 2 (time: 1 or 2) × 2 (initial description: approach or avoid) × 2 (interaction time: thirty seconds or instructions only) mixed-model analysis of variance, in which time was a within-subjects variable, and initial description and interaction time were between-subjects variables. This analysis yielded only one effect of interest: a 2-way interaction between time and order that was unqualified by interaction time, $F(1, 95)=20.056, p<.001$ (see Figure 9).

To assess whether participants rapidly formed implicit responses toward *wugs*, I compared participants' implicit attitudes at time 1. This analysis revealed that participants who played (or read about) the positive round first exhibited significantly more implicit positivity towards *wugs* than those who played the negative round first, $F(1, 95)=18.291, p<.001$. Moreover, there was also strong evidence that participants rapidly revised their implicit attitudes as the meaning of the *wugs* shifted from one round to the next, as evidenced by significant effects of time both for those who initially played or read about the approach round, $F(1, 95)=9.843, p<.01$, and for those who initially played or read about the avoid round, $F(1, 95)=10.241, p<.01$. However, a comparison of implicit preferences at time 2 failed to reveal evidence of a complete reversal of participants' attitudes, $F(1,95)=1.107, ns$.




Figure 9. The two-way interaction between time and round in Study 11. This effect was unmoderated by amount of exposure to *wugs*.

Discussion

In the context of a completely different domain and a different type of relevance, participants' implicit preferences were once again found to shift very quickly, in line with one's most recent experience with an attitude object. Moreover, participants even exhibited these strong shifts in their implicit preferences under conditions in which their exposure to the novel stimulus was exceedingly minimal. Indeed, participants in the *instructions only* condition merely *expected* to interact with *wugs* at a later point in the experiment—essentially abstractly supposing that *wugs* were either good or bad—without ever actually experiencing a positive or negative outcome as a result of

interacting with them. Yet, this expectation appears to be enough for implicit responses to be responsive to recent learning and experience.

Of course, one limitation of these previous studies is that they were conducted exclusively on Cornell undergraduates. To ensure that the consistent effects observed across these studies were not idiosyncratic to the sample population I used in those studies, Study 11 was a conceptual replication of the *instructions only* condition of Study 10 that made use of an online sample that drew upon participants from all over the world.

Study 11: Wugs II

Participants

One hundred and ninety-five participants on Mechanical Turk (<http://www.mturk.com>) completed an online experiment on “video game distractions.” Twenty-four of these participants either: (a) had no variance in their answers to both AMPs (described below; 3 participants), or (b) made excessive errors (>10%) on one or both AMPs (21 participants), and were thus excluded from analyses.

Procedure

The procedures were identical to the previous experiment except for the following changes. First, the entire experiment was administered via a website hosted on <http://cornellpsych.org>. Second, having established that participants’ implicit preferences were equally sensitive to recent experience whether they interacted briefly with the object or merely expected to interact with the object

later in the experiment, only the *instructions only* condition was included in the design. Finally, several changes were made to the instructions of the game to make the descriptions more concrete and engaging. To give participants a helpful mnemonic for remembering the role of the *wugs* in each round of the game, they were told, in the positive context, that they should try to “hug a wug” by touching the *wugs* with their character. If they could do so, the *wug* would give the participant’s character 500 points. In the negative round, participants were told that the *wugs* should be avoided at all costs and that if the *wugs* could catch the player’s character, they would be “mugged by a wug” and they would lose all of their accumulated points.

Results

To assess whether participants’ implicit preferences were sensitive to recent experience, I once again performed a 2 (time: 1 or 2) × 2 (initial description: approach or avoid) mixed-model analysis of variance (ANOVA). This analysis yielded the predicted 2-way interaction, $F(1, 169)=17.957, p<.001$ (see Figure 10).

To assess whether participants quickly formed an implicit preference towards *wugs*, I performed an analysis of participants’ implicit attitudes at time 1. This analysis revealed that participants in the positive description exhibited significantly more positivity toward *wugs* than those in the negative description, $F(1,169)=9.606, p<.01$. To assess whether participants’ implicit preferences were quickly revised in light of recent experience, I assessed the effects of time

separately for each group of participants. This analysis revealed significant shifts in participants' responses from time 1 to time 2, both for participants who initially read the approach description, $F(1,169)=7.775, p<.01$, and for those who initially read the avoid description, $F(1,169)=10.294, p<.01$. Moreover, an analysis of implicit attitudes at time 2 also revealed that participants who read the approach description showed marginally greater implicit positivity towards *wugs* than those who read the avoid description, $F(1,169)=3.585, p=.06$, indicating a full reversal of participants' time 1 preferences.

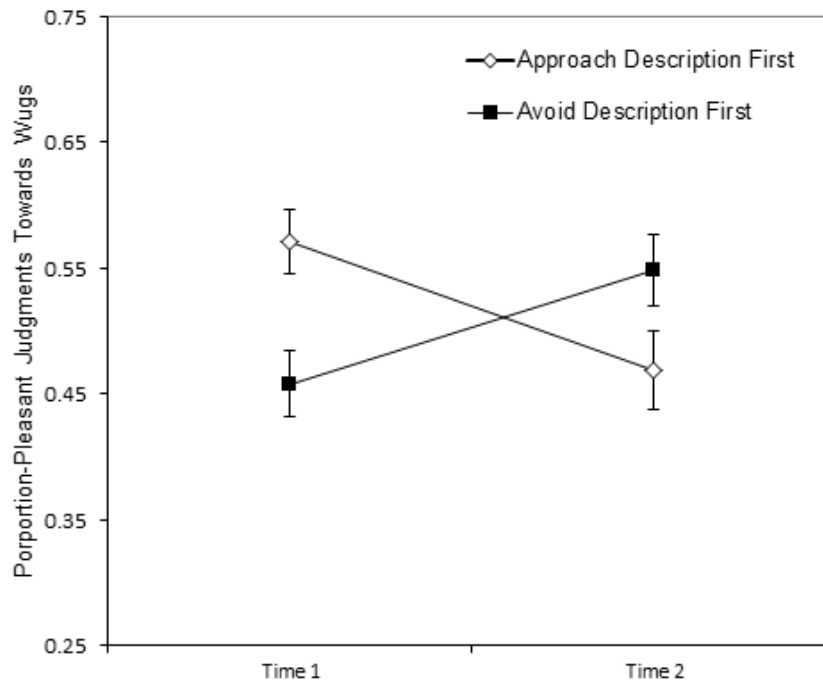


Figure 10. Participants' proportion-pleasant judgments at time 1 and time 2 in Study 12.

Discussion

This study is thus entirely consistent with the results of the previous studies, suggesting that the earlier results were not the product of some idiosyncratic feature of the sample population on which they were conducted. This result thus strongly supports the contention that implicit attitudes can very quickly incorporate recent experience, at least in the context of these two paradigms.

However, one obvious criticism of the previous 5 studies is that, by virtue of using novel attitude objects with which participants had had no previous exposure, the attitudes that participants revised in light of new information were based on very little information and had very little time to be rehearsed or to become crystallized. Although using novel stimuli of this sort has the advantage of controlling participants' previous exposure and history with the attitude objects, thus providing a cleaner test of attitude formation and change, one might argue that such a procedure creates precisely the type of implicit attitude that should be most malleable. Indeed, although previous evidence suggests that even newly-formed attitudes (indeed, even those formed using similar procedures to those used in the previous studies) are insensitive to new, countervailing information (see Gregg et al., 2006, Studies 3 and 4), it remains an open question whether better-learned, more heavily-entrenched attitudes are similarly susceptible to rapid revision in the way that I have shown here.

One reason to remain skeptical, perhaps, of the possibility of rapid revision of more well-established implicit attitudes is that previous research has

suggested that a great deal of counter-learning may be necessary to change an implicit response once it has become well-rehearsed. Rydell and colleagues (2006, 2007), for example, showed that when participants were given extensive, consistent exposure to a target individual named Bob, participants formed an implicit attitude towards Bob that was consistent with their previous exposure. However, once these responses had been established, it took at least 40 trials of counter-learning before participants' implicit responses appeared to exhibit any substantive changes, suggesting that they were relatively unresponsive to new information by virtue of having been heavily-entrenched in the earlier learning paradigm.

However, it is still unclear whether heavily-entrenched implicit attitudes are always insensitive to new information, or if there are certain types of information that can cause rapid revision even after crystallization has taken place. The counter-learning about the target individual in Rydell and colleagues' studies, for example, involved mildly negative behaviours like parking illegally or missing an appointment with a friend. Are there more dramatic types of revelations about an individual that can have a substantive impact on implicit attitudes? That is, if participants learn information that is less like learning that an individual committed a minor infraction and more like learning that one's spouse has been having an extra-marital for the previous several years, can it exert a greater influence on participants' implicit responses?

This was the goal of the final study. Participants first learned about a novel target named Bob, who was portrayed as universally positive over a lengthy learning paradigm. Then, after assessing participants' implicit attitudes towards Bob, they were given just one additional piece of information about him: that he had recently been convicted of molesting children (see Castelli et al., 2004). The primary question of interest was whether participants' well-rehearsed positivity towards Bob could be undone by this single piece of dramatic, extremely negative information.

Study 12: Who is Bob, Really?

Method

Participants. Fifty-eight Cornell undergraduates participated in exchange for credit or \$5. Two participants were excluded due to a lack of variance in their responses on the AMP, and one participant was excluded due to excessive errors on both AMPs (>10%). This left a final sample of 54 participants.

Procedure. Participants first completed a learning paradigm in which they read behavioral statements about a target individual named Bob (e.g., "Bob gave a hitchhiker a ride to a shelter"), identified by a picture presented directly above the behavioral statement.⁷ Their task was to assess whether they thought each target behavior was or was not characteristic of Bob. After reading each behaviour, participants responded by pressing the *c* key (characteristic) or *u* key

⁷ Following a similar protocol to Rydell and colleagues' (2006) procedures, one of 6 different pictures of white males (judged to be of similar physical attractiveness in a pre-test) was used to represent Bob, randomly assigned for each participant. Whenever a particular target served as Bob, the other 5 targets served as neutral stimuli in the implicit measure. There were no effects of target picture in any of the analyses, and this manipulation is not discussed further.

(uncharacteristic). They then received immediate feedback about their guess, which consisted of one word displayed in the center of the screen: either the word “correct” printed in blue text, or “incorrect” printed in red text, followed below it by a summary of the meaning of the feedback (e.g., “Giving a hitchhiker a ride to a shelter is characteristic of Bob”). To ensure that participants initially exhibited strong positivity towards Bob, 100% of the positive behavioral statements were said to be characteristic of Bob, while 100% of the negative statements were said to be uncharacteristic of him. The order of the behavioral statements was randomly determined by the computer.

Assessing implicit attitudes. After the 100 learning trials, participants’ implicit attitudes towards Bob were assessed using an AMP. In this case, the primes consisted of 30 pictures of Bob, as well as 30 pictures of 5 unfamiliar targets (that is, 6 pictures each).

Changing impressions of Bob. To assess the effects of new, countervailing information on participants’ now-well-rehearsed implicit reactions to Bob, participants were told that they would learn one new piece of information about him. All participants were told that they should pay close attention because this new piece of information was more recent and could potentially differ from the impression they had otherwise formed of him. In the *experimental* condition, participants read the statement, “Bob was recently convicted of molesting children.” In the *control* condition, participants read the statement, “Bob recently bought a soda.” In both cases, participants were reminded that this

piece of information was characteristic of Bob and that they would be tested for their memory of this statement later in the experiment.

Participants then completed a second AMP, which followed the same protocol as the first. They then completed a short questionnaire assessing their explicit attitudes towards Bob by judging his likeability from 1 (very unlikeable) to 9 (very likeable) as well as how *bad—good, mean—pleasant, disagreeable—agreeable, uncaring—caring, and cruel—kind* they considered him to be, all on 9-point likert scales.

Results

To assess changes in participants' implicit attitudes towards Bob, I performed a 2 (time: 1 or 2) × 2 (new information: child molester or control) mixed-model analysis of variance with time as a within-subjects factor and new information as a between-subjects factor. This analysis yielded the predicted 2-way interaction, $F(1, 53)=10.327, p<.01$ (see Figure 11).




Figure 11. Participants’ proportion-pleasant judgments towards Bob before (time 1) and after (time 2) learning either that he was convicted of molesting children (experimental) or that he bought a soda (control) in Study 12.

To assess whether participants had initially developed implicit positivity towards Bob, I performed a single-sample *t*-test on participants’ implicit responses at time 1, comparing them against the value .5, which represents a neutral implicit attitudes. This analysis revealed a marginally significant difference in the positive direction, indicating that participants had indeed formed relatively positive implicit attitudes towards Bob after learning extensively about his many positive attributes, $t(54)=1.959, p=.055$.

Next, to assess whether participants’ well-rehearsed implicit responses towards Bob were sensitive to the new information that they learned about him,

I examined the effects of time separately for participants in the *control* and *experimental* conditions. This analysis revealed a significant shift in participants' attitudes when they learned that Bob was convicted of child molestation, $F(1,53)=10.096, p<.01$, but no significant changes in implicit preferences when participants were exposed to neutral information at time 2, $F(1,53)=1.914, ns$. These differential shifts resulted in a significant difference between control and experimental participants' attitudes at time 2, $F(1,53)=17.894, p<.001$. Moreover, when participants' implicit attitudes at time 2 were compared, using a single-sample *t*-test, against the neutral value of .5, there was evidence that participants in the experimental condition had exhibited a full reversal in their responses, showing significant negativity towards Bob after learning of his conviction, $t(27)=2.671, p<.05$, whereas control participants retained their significant implicit positivity towards him, $t(26)=3.282, p<.01$.

This result thus suggests that even a well-rehearsed implicit attitude can be sensitive to new information that is strongly inconsistent with previous experiences with an attitude object (cf. Rydell et al., 2006; Rydell et al., 2007). Indeed, even just a single piece of propositional information about Bob—that he had been convicted of a serious crime—was enough for participants to exhibit a complete reversal from all of the information that they had previously acquired about him in the previous 100 trials.

General Discussion

Attitudes theorists have claimed that implicit attitudes are, by virtue of their reliance on slow-learning associative mechanisms, rather slow to develop, and rather insensitive to recent, relevant experiences with an attitude object once established. Yet, in the 6 studies reported here, participants very quickly formed an implicit attitude towards a novel stimulus with minimal learning and experience—findings that are inconsistent with traditional assumptions about the learning characteristics of implicit attitudes (Bassili & Brown, 2005; Conrey & Smith, 2007; Rydell & McConnell, 2006; Smith & DeCoster, 2000; Strack & Deutsch, 2004; Wilson, Lindsey, & Schooler, 2000; cf. Gawronski & Bodenhausen, 2006, 2011), but that join emerging evidence of rapid formation (e.g., Ashburn-Nardo et al., 2001; Castelli et al., 2004; Gregg et al., 2006).

More importantly, participants' implicit attitudes were also found to be susceptible to quick revision in light of new, countervailing information (cf. Gawronski & Strack, 2004; Gregg et al., 2006; Wilson et al., 2000): in studies 8, 9a and 9b, participants' implicit attitudes towards members of novel groups exhibited significant changes immediately after learning that their original group assignment feedback was the result of a computer error; in Studies 10 and 11, participants' implicit attitudes towards a novel attitude object exhibited changes in line with manipulations of the meaning of the novel object; and in Study 12, participants' implicit responses to a target individual were dramatically affected by a single piece of propositional learning. To my knowledge, this constitutes the first evidence for rapid revision of implicit attitudes and challenges claims about

the nature of the processes underlying implicit attitude change, particularly theories predicting slow and gradual learning.

Why did I find evidence for rapid revision where previous work did not? I propose that an important difference between the present studies and earlier work is that the novel objects were simply made self-relevant and important to participants' experiences in these experiments. In Studies 8, 9a and 9b, participants were not exposed to hypothetical groups in vignettes, but rather to an *ingroup*—that is, one for which a change of ingroup assignment is more meaningful (e.g., Campbell & Tesser, 1985; Tesser, 1988). Similarly, in Studies 10 and 11, the manipulation of the meaning of the attitude object with which participants interacted was directly relevant to the goal of doing well in an upcoming task. Finally, in Study 12, the piece of information that participants learned about the individual was particularly diagnostic of him, perhaps casting his previous behaviours in a different light, and evoking a sense of deceit and betrayal. Of course, it is important to note that because I did not directly manipulate self-relevance in the current studies, I cannot be certain of this interpretation of the results, and additional research will be required to discern the conditions under which implicit attitudes are or are not sensitive to revision once established.

These results are also inconsistent with previous studies that have employed similar learning paradigms to create well-rehearsed implicit responses (e.g., Rydell & McConnell, 2006; Rydell et al., 2007). In this earlier work,

researchers found that implicit attitudes formed rather slowly and inexorably as additional information about the target was acquired, and changed equally slowly and only with extensive countervailing information. Why did I find evidence of implicit attitude change after a single trial of propositional learning while in their studies they found that it took as much as 40 pieces of counterattitudinal information to cause implicit attitudes to exhibit any substantive changes? Although currently I can only speculate, there are a number of possibilities. First, it is clear that not every piece of information that one learns about a target is equally diagnostic. Learning that someone has an overdue parking ticket is substantively different from learning that an individual has molested children, and it is encouraging that the latter behaviour can be quickly incorporated into one's implicit responses, even if the former cannot.

Second, it may be the case that the nature of the information that was presented here is such that it is not only viewed as a single instance of exceedingly negative behaviour, but also as a behaviour that serves to cast all of the previous information about the target in a different light. What was previously seen as an exceedingly charitable and generous individual suddenly becomes a self-serving, overly compensating criminal. In this way, a single piece of propositional information served to undermine all of one's previous learning about the target, which may have made it particularly effective in changing participants' implicit responses.

Finally, being convicted of a serious crime carries with it a great deal of evaluative baggage—that is, a large amount of exceedingly negative evaluative information about the group “*child molesters*” that may be easily imported into people’s implicit responses towards a specific exemplar (Castelli et al., 2004). Indeed, one reason that the results of the minimal group studies may have been so effective in changing participants’ attitudes is by virtue of their association with the self—a construct that carries with it a large number of evaluative associations that can, perhaps, be readily generalized to objects that become associated with the self (see also, Gawronski, Bodenhausen, & Becker, 2007). In the same way, attitudes towards Bob may have been quickly associated with the group “*child molesters*”, and individuals’ prior evaluative knowledge about this group could then be readily applied to him.

Although teasing apart these explanations and the potential moderators of these effects will require additional research, these findings nonetheless provide an existence proof that implicit attitudes can possess learning and re-learning characteristics that are inconsistent with those that are often thought to characterize implicit memory processes in the social cognition literature (see Amodio & Ratner, 2011). Whereas models of implicit attitudes have generally assumed that associative processes afford only a single learning trajectory, cognitive and social psychologists and neuroscientists have argued that associative processes (sometimes called implicit memory) in fact consist of multiple, interacting systems and processes (e.g., Amodio & Ratner, 2011;

Poldrack & Foerde, 2007; Squire & Zola, 1996). The current results thus suggest that implicit attitudes may rely on more than one type of associative or implicit process, and may thus exhibit different learning characteristics under different circumstances. This view is also consistent with recent work showing that implicit measures (e.g., of attitudes) do not reveal a single process, but instead inevitably capture multiple processes (e.g., Conrey, Sherman, Gawronski, Hugenberg, & Groom, 2005), some of which may allow for fast learning.

I have claimed that I consider this to be the first evidence of rapid revision of implicit attitudes. However, one might suggest that there is ample evidence for rapid implicit attitude change already in the literature (e.g., Gawronski & Bodenhausen, 2006, 2011). Indeed, research has shown that implicit attitudes towards the same stimulus are heavily context-dependent, responding to current emotional states, accessible goals, and a variety of other contextual cues (for a review, see Gawronski & Sritharan, 2010). For example, Dasgupta and Greenwald (2001) reported that people showed reduced implicit negativity toward African-Americans when they had recently viewed photographs of liked (e.g., Michael Jordan) versus disliked (e.g., Mike Tyson) exemplars. This does not necessarily indicate, however, that people's implicit attitudes toward the group have *changed* in the way that I mean here (e.g., see Fazio, 2007; Ferguson & Fukukura, 2012). Rather, such results may indicate the activation of different aspects of the same (unchanged) mental representation at the time of implicit evaluation. And, most importantly, such context-specific implicit attitudes may

still require extensive experience and develop only gradually over time. In this way, context dependence of implicit attitudes does not speak to how easily implicit attitudes can be revised in line with new information. The current studies instead examined novel attitude objects with which participants cannot have had previous experience, and thus provide a more sensitive and controlled test of rapid formation and revision at the level of the mental representation.

The current studies highlight the possibility that implicit attitudes—even those based on a relatively large amount of experience—may have a greater capacity for change than current theories assume. The implication, then, is that implicit attitudes need not be a source of bias in situations in which recent, relevant experiences should play an important role in people’s behaviour, at least under the circumstances captured in the studies reported here. Indeed, theories predicting slow and gradual learning may very well be in need of (rapid) revision.

CHAPTER V: General Discussion

I.

What was the source of Timothy Cook's intuition on that day in March of 1998 that left him with the feeling that Steve Jobs had charted a course that would dramatically change the future of a failing company in disarray? We will likely never know exactly what it was that gave rise to that stunningly prescient insight, nor will we likely ever know what it was that made his intuition seem so compelling that he chose to side with that feeling rather than all of the various logical considerations that had suggested it would be a mistake. However, in the preceding chapters, I have tried to advance our understanding of situations like Cook's, exploring the general principles that might govern how people come to trust one decision strategy over another and whether they are likely to be well-served by doing so.

In Chapter II, I sought to explore the descriptive question of how people can be influenced to trust the dictates of intuition or reason—that is, the predictors of trusting intuitions or ignoring them in favour of rational analyses. I found consistent evidence to support the contention that one way people attempt to resolve such conflicts is by looking both outward—that is, to properties of the decision and the context in which it occurs—as well as inward—that is, to properties of their current internal mental or bodily state. I found that when aspects of a person's internal state closely resemble the features of intuitive or rational judgment, this can serve as a diagnostic cue that

leads them to trust one source of input over the other. These studies that suggest that there are indeed systematic influences on the extent to which an intuitive judgment “feels right,” and that these influences can affect the choice strategy that people adopt in their decision-making.

In Chapter III, I sought to turn my attention to the prescriptive question of whether intuitions can serve as a reliable guide for successful decisions. In this chapter, I honed in on an assumption that has been made about the source of gut feelings, which is that they can implicitly capture the breadth of one’s prior experiences with an object, recording without awareness or intention many features such as frequency and value (e.g., Betsch et al., 2001; Betsch, 2008). However, such information is only useful, I argued, to the extent that the processes that govern the acquisition of information from the environment possess a necessary filter that can cast aside fallacious or irrelevant experience. Ultimately, this hypothesis turned out to be untestable due to consistent failures to replicate previous research using a well-known paradigm in the judgment literature, rendering this an open question for future research, perhaps employing other more reliable paradigms from the judgment and decision-making literature.

The purpose of Chapter IV was to extend the prescriptive analysis of implicit and explicit processes to the domain of implicit evaluation. In this chapter, I honed in on another assumption about the nature of implicit processes which is that they are relatively insensitive to recent relevant experiences, thus

making them a source of bias in situations in which relevant experience should serve as an appropriate, normative guide for behaviour. I found that there is reason to believe, however, that such an assumption may be unwarranted and that implicit attitudes can indeed come to rapidly reflect recent experience with an attitude object, particularly in situations in which the target is relevant to the individual and has important consequences for their future behaviour. I found that when these minimal conditions of relevance were met, implicit attitudes shifted rather quickly and rather markedly in line with new, countervailing information. Moreover, this occurred even for heavily-entrenched implicit attitudes that had been formed through extensive and consistent repetition over a lengthy learning paradigm. This result thus suggests that implicit attitudes need not necessarily serve as a source of bias when recent experiences are inconsistent with past learning.

With these findings in mind, I now take a broader view of the results across these three lines of work, exploring some of their implications and some potential avenues for future research.

II. When do implicit processes ignore irrelevant past experiences and when do they incorporate them?

There is a potential inconsistency in the theoretical arguments I made between Chapter III and Chapter IV. In Chapter III, I proposed that theoretical conceptions of implicit processes might imply that intuitive judgment lacks the ability to filter unnecessary or irrelevant experience, and thus rather

promiscuously incorporates all aspects of one's prior exposure with an attitude object irrespective of its truth value. Cast in a different light, it might appear as though this claim implies that intuitions lack any mechanism by which to discount elements of past experience. However, in Chapter IV, I argued that implicit attitudes can come to reflect one's most recent experiences with an attitude object, suggesting that implicit attitudes can indeed discount past experiences in favour of only a subset of one's prior exposure with the attitude object. How are we to reconcile this apparent inconsistency?

As I argued, an essential ingredient in whether rapid learning and revision takes place seems to be the extent to which stimuli are relevant and important for future outcomes. However, in stark contrast to the manipulations of relevance that I employed in the studies in Chapter IV, the paradigms on which the claim that intuitions can capture the breadth of one's past experiences are generally designed in such a way as to ensure that the stimuli are relatively unimportant to participants. Indeed, Betsch and colleagues (2001) have proposed that one of the requirements necessary for intuitions to reflect a summary evaluation of the value of a decision object is that individuals are mildly distracted during the information acquisition process. Participants are told that their primary focus during the experiment should be on the print ads, and the stock information is described to them as merely a tool for distracting them while they attempt to memorize other information. In this way, the stock information has very little relevance or importance, and these are precisely the

situations in which we would expect that intuitions fail to discount past experiences in favour of more recent, relevant ones.

I hasten to note, however, that it is still unclear whether intuitions promiscuously capture all of one's past experiences by virtue of being unable to submit this claim to empirical test. Moreover, even if intuitions had been shown to be influenced by fallacious or irrelevant information, it is still unclear as to whether such an effect would still occur even in a situation in which the stimuli were made more self-relevant—say, by indicating to participants that their compensation would be based on stock performance.

III. The implicit: One system/process or many?

In Chapter IV, I made the argument that implicit memory—the mechanism by which implicit evaluation is thought to operate—has been characterized by cognitive psychologists and neuroscientists not as a single, monolithic entity, but rather as several interacting memory systems that each possess different properties, learning characteristics, and developmental trajectories (Amodio & Ratner, 2010; Poldrack & Foerde, 2007; Squire & Zola, 1996). This is an important observation because it suggests that if theorists characterize the implicit system (or process) as a single entity, they may potentially gloss over important differences in the way implicit processes contribute to judgment and behaviour under different circumstances. A number of researchers have proposed that we might profitably make further distinctions within the broad categories of implicit and explicit in order to capture distinct

abilities with different kinds of properties (Evans, 2008; Glockner & Witteman, 2010; Gore & Sadler-Smith, 2011; Hogarth, 2010; Stanovich, 2004, 2011).

Stanovich (2004, 2011), for example, has characterized System 1 not as a unique system with undifferentiated properties but as The Associative *Set of Systems* (TASS; emphasis added) to highlight the notion that it likely encompasses many different abilities that exert greater influence under different circumstances.

In the context of intuitive judgment, researchers have often lamented the lack of conceptual clarity regarding the study of intuition, failing to properly define what intuition is and what it is not (Betsch, 2008; Glockner & Witteman, 2010; Hammond et al., 1987; see also Shah & Oppenheimer, 2008, for a similar concern about the related concept of “heuristic”). More recently, researchers have sought resolve this lack of clarity by conceptualizing intuition not as a single type of judgment, but rather, not unlike contemporary conceptions of implicit memory, as a collection of interacting implicit processes that differentially operate in different types of judgments. In a recent review, for example, Gore & Sadler-Smith (2011) sought to “unpack intuition,” proposing that it might be best characterized as several distinct types of domain-general mechanisms such as heuristic processing and associative learning, as well as several distinct types of outputs, including problem-solving intuitions (which are thought to form the basis of intuitive expertise), social intuitions, creative intuitions, and moral intuitions. Glockner & Witteman (2010) have, in a somewhat different conception, proposed 4 separate sub-domains of intuition that may differ in

terms of their effects on a host of different variables (e.g., physiological arousal, decision time, etc.)

Although it is still unclear where the proper demarcations may lie, such distinctions may serve to resolve conceptual confusion about how intuition should be properly defined and when intuitive judgment is likely to be beneficial. Indeed, one potential resolution of the observation that researchers in different traditions (e.g., heuristics and biases vs. naturalistic decision-making) have come to quite different conclusions about the value of intuitive judgment is that these researchers may have unwittingly emphasized different types of intuiting—some of which are beneficial and others of which are more likely to be a source of bias (see Kahneman & Klein, 2009, for a more detailed discussion of the distinctions between these traditions and their possible synthesis).

III. Do people have insight into when implicit processes are beneficial?

An interesting implication of the task and state cuing hypotheses (discussed in Chapter II) is that there is a great deal of variation in the extent to which people come to trust their intuitive judgments or not. People appear to have lay theories about the relative quality of intuitive decision-making under different circumstances, trusting it in some circumstances and not in others. An important theoretical question thus concerns whether people's lay theories about when intuition is helpful align with the situations that have been identified by researchers as ones in which intuitive judgment is relatively more reliable. Are people's intuitions about their intuitions correct? Do they have any accurate

insight into when their intuitions are surer guides to sound decisions versus when they are more often a source of error?

As the failures to replicate in Chapter III highlight, we still know very little about when intuitive processes are beneficial or not. Thus, the answers to such questions are still perhaps currently out of reach. However, there are a few broad conclusions that warrant further exploration. One important observation, for example, is that some of the factors that served to pull participants in the direction of their rational analyses are precisely the situations in which researchers have proposed that rational analyses are likely to be most unreliable—that is, in situations in which participants encounter difficulties in their decision-making, either because the problem is complex or the stimuli are difficult to process (Betsch et al., 2001; Dijksterhuis, 2004; Dijksterhuis et al., 2006). If these researchers are correct in their assertion that intuitions have an edge on rational analyses in such situations, it suggests that people lack any meaningful insight into this phenomenon. Indeed, Dijksterhuis (2004), in his discussion of his Unconscious Thought Theory, notably described the assertion that people are better-served by intuitions as problems complexity increases as a “somewhat counterintuitive idea” (p. 597), suggesting that it may defy many people’s intuitions about their intuitions.

It is also worth noting that the manipulations of participants’ mental states at the time of their decisions that I reported in Chapter II were entirely incidental to the actual decision, and normatively should perhaps not have had

any strong influence on the choices that people made. This suggests that participants' lay theories of the power of intuition can indeed break down, at least under some well-controlled, experimentally-induced circumstances. Moreover, because of the nature of the paradigm used in those studies, any mindset that served to influence people to trust their intuitions simultaneously served to reduce their likelihood of receiving additional compensation, making it a normatively inferior choice strategy. The fact that even in such situations people can still often be relatively more persuaded by their intuitive judgments than by a reasoned analysis is thus perhaps rather telling.

More generally, there is a strong theoretical precedent to suggest that individuals are largely ignorant of the experiences or processes that serve as the antecedents of their gut feelings. Many researchers have proposed that a defining feature of intuitions is that their source (and, by extension, any biases that may exist in that source) is wholly unavailable to conscious introspection. Hogarth (2010) has proposed, for example, that individuals lack meta-cognitive insight into whether their intuitions have been honed in a *kind* or a *wicked* environment, thus leading them to place their trust in their intuitive judgments even when they are unreliable.

Nonetheless, there is still work to be done in determining whether people have any accurate insight into the quality of their intuitive decision-making—that is, work that examines whether the variance that is observed in people's likelihood of trusting intuitions bears some resemblance to the likelihood of

intuition leading to good decisions. Perhaps there are indeed reliable cues that exist in the environment that serve as an indicator to individuals that their intuition may be faulty. If such cues exist, are people able to pick up on them, either consciously or non-consciously?

To take an example, consider Hogarth's central claim that the quality of intuitive decision-making is a product of the extent to which the environment in which one makes a decision is representative of the environment in which learning takes place. One cue that may exist that could serve as an indicator of the quality of one's intuitive insights, then, is if there have been any changes in the structure of the environment from the time of learning to the time of making a later decision based on an intuitive judgment. Are people sensitive to such changes in the environment? And could such changes serve as a cue to individuals that their intuitions may be relatively less useful and reliable? Suppose, for example, that participants are given relevant experience in one environment identifiable by a particular cue—say, a particular background colour. Are participants more likely to side with their intuitions when making a later judgment if the same cues are present than they are if these cues differ? This is a question that awaits future research.

IV. Future directions

Implicit Attitudes

In Chapter IV, I argued that implicit attitudes may reflect recently-learned information about an attitude object, even in situations in which one has had

extensive and consistent previous exposure with the object—a result that is inconsistent with traditional and contemporary theoretical conceptions (e.g., Bassili & Brown, 2005; Conrey & Smith, 2007; Rydell & McConnell, 2006; Smith & DeCoster, 2000; Strack & Deutsch, 2004; Wilson, Lindsey, & Schooler, 2000), and that runs counter to a number of findings already in the literature (Rydell & McConnell, 2006; Rydell et al., 2006; Rydell et al., 2007; McConnell et al., 2008).

An additional assumption that has been made by attitudes theorists with respect to the learning characteristics of implicit attitudes is that they are especially responsive to affectively-laden associative information and especially insensitive to propositional knowledge. McConnell and colleagues (2008), for example, have shown that implicit attitudes are, in the context of person perception, especially sensitive to affectively-laden group association cues, such as whether a person is obese or especially attractive. They have suggested, moreover, that cues of this sort can swamp other aspects of one's experience, even if that experience consists largely of instances of behaviour that are inconsistent with the affective cue. For example, when participants learned about a target named Bob who was especially obese, McConnell and colleagues found that participants' implicit attitudes were entirely driven by Bob's obesity, and rather insensitive to whether Bob engaged in positive or negative behaviours. An interesting question for future research, then, is whether the effects I observed in Study 12 are replicable in a context in which such affective information is salient. Does learning that an especially attractive Bob has

recently molested children have a similar impact on one's implicit responses towards him as it does when he is of average attractiveness? Or is his high attractiveness enough to prevent rapid revision in such cases?

Another theoretically meaningful question for future research concerns whether valence asymmetries may exist in the learning and re-learning of relevant evaluative information. In Study 12, for example, participants were initially given positive information about Bob that was later reversed with a single instance of exceedingly negative propositional information about him. However, it remains an open question whether rapid revision can occur for similarly well-rehearsed negative implicit attitudes. Although the studies that made use of the *Wugs* video game (Studies 10 and 11) failed to uncover any evidence of such asymmetries—that is, equivalent effects were observed for those whose initial experience was positive or negative—there is a strong theoretical precedent for such effects in affective learning. Rozin & Royzman (2001) observe, for example, that an entire plate of pasta can be contaminated by a single cockroach, but that no amount of pasta can make a plate of cockroaches edible. In the same way, perhaps, at the implicit level, no number of kind acts can make up for consistent moral failures.

Another interesting question for future research concerns the generalizability of implicit attitudes to other novel targets. Research has suggested that affective learning about one target can be generalized to other targets that bear a perceptual resemblance to it (Verosky & Todorov, 2010). An

interesting question in light of the rapid revision observed across studies 8-12 is the extent to which the learning of new, countervailing information about a target can similarly generalize, at the implicit level, to other targets.

For example, consider one's implicit evaluations of one's romantic partner. What happens to these implicit evaluations when one learns that one's partner has cheated on him or her? The research reported here would suggest that such devastating information may have an immediate effect on people's (presumably well-rehearsed) implicit responses towards their partner. But what happens to one's implicit evaluations of targets that resemble one's partner? Do such evaluations similarly shift in line with this recently-learned information? Even more interestingly, suppose that a novel target that resembles one's partner has recently been convicted of molesting children. Does learning this information have any substantive impact on one's implicit responses to one's partner, by virtue of the resemblance between the two?

Intuitions about Intuitions

The focus of Chapter II was on the meta-cognitive task that individuals face when they are pulled in different directions by implicit and explicit processes. In these studies, I focused exclusively on hypothetical intuition/reason conflicts in which the conflict was made explicit to participants, and other well-studied experimental paradigms that have sought to capture situations in which intuitive and rational processes pull in opposing directions.

However, as I outlined earlier in this chapter, researchers have increasingly begun to recognize that intuition may consist of many different types of implicit processes—including problem-solving intuition, social intuition, and moral intuition—that operate under different circumstances and possess different antecedents, properties and learning characteristics. One of these sub-domains concerns creative intuitions—the flashes of insight that people experience when they attempt to solve a creative problem or generate a new idea. This is, of course, an area of long-standing interest to cognitive and social psychologists, and a great deal of progress has been made in understanding the antecedents of creative intuition. However, like the analysis I performed in Chapter II, it is similarly true that considerably less is known about people’s intuitions about their creative intuitions—that is, the meta-cognitive problem that individuals face as they attempt to judge whether an idea that they have had is worthwhile and potentially fruitful, or flawed, trite, or otherwise uninspired.

The analysis in Chapter II suggests that there may be interesting situational factors that could influence this meta-cognitive task, indicating, perhaps, that judgment of the value of creative ideas might be influenced by aspects of the type of task on which the creative ideas focus or on properties of an individual’s mental state at the time that an insight occurs. Moreover, people’s intuitions about their creative intuitions may differ in important ways from their actual quality.

The Nature of Intuitions

In the same way that researchers in social cognition have made strong claims about the developmental trajectories of implicit attitudes, so too have researchers in judgment made strong claims about the speed with which intuitions develop (e.g., Plessner, Betsch, & Betsch, 2008). A number of researchers operating within different theoretical frameworks have emphasized the slow-developing nature of intuitive impulses and their relative stability once they become established through extensive experience. Simon (1987) famously contended, for example, that “intuition and judgment are merely experiences *frozen into habits*” (p. 63; emphasis added). Betsch (2008), however, has probably made this assertion most explicitly, claiming that:

“intuition primarily capitalizes on prior knowledge acquired via a slow learning system rather than on recently encountered information kept in short-term memory”

(...)

“herein lies one asset of deliberation: It is open to new evidence enabling the individual to contextualize and to rapidly adapt to changing situations. In contrast, intuition is more likely to produce conservatism in judgment and decisions because it is primarily driven by consolidated knowledge” (page 4)

However, empirical tests of these claims are, to my knowledge, non-existent. This suggests that a fruitful avenue for future work may be to subject intuitive judgments to empirical tests similar to those employed in Chapter IV, Are intuitions relatively more susceptible to rapid shifts in a relatively short period of time than current theoretical conceptions assume? Can they exhibit large reversals on the basis of relatively more recent learning rather than producing “conservatism” driven by “consolidated knowledge”?

Relatedly, as I discussed in Chapter IV, social cognition researchers have, in recent years, begun to recognize that implicit attitudes are heavily influenced by the context in which they are activated, responding to current emotional states, accessible goals, currently activated mental representations, and a host of other factors (for a review, see Gawronski & Sritharan, 2010). As I argued at the time, such results do not necessarily speak to the ability of implicit attitudes to change, but they do nonetheless speak quite strongly to the temporal (in)stability of implicit responses, suggesting that subtle changes in context can give rise to widely varying implicit evaluations. Are intuitive judgments similarly susceptible to contextual influences? Hogarth (2002) has suggested, for example, that intuitions may exhibit what he called “field” or “set” effects, in which intuitive judgments can be influenced by the state of an individual at the time that a stimulus is encountered. However, taking a broader view, exactly what role does context play in the output of an intuitive process?

V. Concluding Remarks

John Bargh and colleagues have characterized the shifts in social psychologists’ views over the last thirty years of the role of implicit and explicit processes in guiding behaviour as a “sea-change” (Bargh & Huang, 2009, p. 127). The phrase is quite apt. We are at a point in the history of social psychology in which the role of non-conscious processes in guiding aspects of judgment and behaviour is taken more seriously than perhaps it ever has been. Certainly, there are still many researchers that emphasize the virtues of clear, rational thought,

but the other side of the debate now has many more advocates than it had in previous decades.

Like any complex issue, the truth is likely to lie somewhere in between these views and future research will be faced with the task of understanding when and why non-conscious processes are beneficial in how we decide and how we act. In this dissertation, I have tried to make some progress on these issues by exploring the interaction between implicit and explicit processes both descriptively and prescriptively. Although there is still a great deal of work to be done, I can only hope that the recent (relevant?) experience of considering the ideas I have laid out in this dissertation may contribute in some small way to a revision in one's intuitions about intuitive judgment and implicit and explicit processes more generally.

References

- Alter, A. L., & Oppenheimer, D. M. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychology Review, 13*(3), 219-235.
- Alter, A., Oppenheimer, D., Epley, N., & Eyre, R. (2007). Overcoming intuition: Metacognitive difficulty activates analytical reasoning. *Journal of Experimental Psychology: General, 136*, 569–576.
- American Psychiatric Association, 1994. *Diagnostic and Statistical Manual of Mental Disorders, DSM-IV, fourth ed.* American Psychiatric Association, Washington, DC.
- Amodio, D. M. & Devine, P. G. (2006). Stereotyping and evaluation in implicit race bias: Evidence for independent constructs and unique effects on behavior. *Journal of Personality and Social Psychology, 91*, 652-661.
- Amodio, D. M., & Ratner, K. G. (2011). A memory systems model of implicit social cognition. *Current Directions in Psychological Science, 20*, 143-148.
- Ashburn-Nardo, L., Voils, C. I., & Monteith, M. J. (2001). Implicit associations as the seeds of intergroup bias: How easily do they take root? *Journal of Personality and Social Psychology, 81*, 789-799.
- Bargh, J. A., & Ferguson, M. J. (2000). Beyond behaviorism: On the automaticity of higher mental processes. *Psychological Bulletin, 126*, 925-945.
- Bargh, J. A., & Huang, J. Y. (2009). The selfish goal. In G. B. Moskowitz & H. Grant (Eds.), *The Psychology of Goals* (pp. 127-150). New York: Guilford.

- Baron, J. (1988). *Thinking and deciding*. New York, NY: Cambridge University Press.
- Bassili, J. N., & Brown, R. (2005). Implicit and explicit attitudes: Research, challenges and theory. In D. Albarracín, B. T. Johnson, & M. P. Zanna (Eds.), *Handbook of Attitudes and Attitude Change*. Mahwah, NJ: Erlbaum.
- Betsch, T. (2008). The nature of intuition and its neglect in research on judgment and decision-making. In Plessner, C. Betsch, & T. Betsch (Eds.). *Intuition in Judgment and Decision-Making*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Betsch, T., Plessner, H., Schwieren, C., & Gütig, R. (2001). I like it but I don't know why: A value-account approach to implicit attitude formation. *Personality and Social Psychology Bulletin*, *27*, 242-253.
- Bless, H., Bohner, G., Schwarz, N., & Strack, F. (1990). Mood and persuasion: A cognitive response analysis. *Personality and Social Psychology Bulletin*, *16*, 311-345.
- Bless, H., Mackie, D. M., & Schwarz, N. (1992). Mood effects on attitude judgments: Independent effects of mood before and after message elaboration. *Journal of Personality and Social Psychology*, *63*, 585-595.
- Bodenhausen, G. V., Kramer, G. P., & Süsser, K. (1994). Happiness and stereotypic thinking in social judgment. *Journal of Personality and Social Psychology*, *66*, 621-632.

- 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*, 45-62.
- Brunswik, E. (1957). Scope and aspects of the cognitive problem. In H. Gruber, K. R. Hammond, & R. Jessor (Eds.), *Contemporary approaches to cognition* (pp. 5–31). Cambridge, MA: Harvard University Press.
- Cahill, L., & McGaugh, J. L. (1990). Amygdaloid complex lesions differentially affect retention of tasks using appetitive and aversive reinforcement. *Behavioral Neuroscience, 104*, 532-543.
- Campbell, J. D., & Tesser, A. (1985). Self-evaluation maintenance processes in relationships. In S. Duck & D. Perlman (Eds.), *Personal relationships* (Vol. 1). London: Sage.
- Castelli, L., Zogmaister, C., Smith, E. R., & Arcuri, L. (2004). On the automatic evaluation of social exemplars. *Journal of Personality and Social Psychology, 86*(3), 373-387.
- Chabris, C. F., & Simons, D. J. (2010). *The invisible gorilla: And other ways our intuitions deceive us*. New York: Crown.
- Chandler, J. J., & Pronin, E. (in press). Fast Thought Speed Induces Risk Taking. *Psychological Science*.
- Chartrand, T. L., & Bargh, J. A. (2002). Nonconscious motivations: Their activation, operation, and consequences. In A. Tesser & D. Stapel (Eds.),

- Self and motivation: Emerging psychological perspectives* (pp. 13–41).
Washington, DC: American Psychological Association.
- Damasio, A. R. (1994). *Descartes' error: Emotion, reason and the human brain*.
New York, NY: HarperCollins.
- Dane, E., & Pratt, M. G. (2007). Exploring intuition and its role in managerial
decision making. *Academy of Management Review*, *32*, 33–54.
- Dasgupta, N., & Greenwald, A.G. (2001). On the malleability of automatic
attitudes: Combating automatic prejudice with images of admired and
disliked individuals. *Journal of Personality and Social Psychology*, *81*, 800-
814.
- Denes-Raj, V., & Epstein, S. (1994). Conflict between intuitive and rational
processing: When people behave against their better judgment. *Journal
of Personality and Social Psychology*, *66*, 819–829.
- DeSteno, D. A., Dasgupta, N., Bartlett, M. Y., & Caidric, A. (2004). Prejudice from
thin air: The effect of emotion on automatic intergroup
attitudes. *Psychological Science*, *15*, 319-324.
- Dijksterhuis, A. (2004). Think different: The merits of unconscious thought in
preference development and decision making. *Journal of Personality and
Social Psychology*, *87*, 586–598.
- Dijksterhuis, A., Bos, M. W., Nordgren, L. F., & van Baaren, R. B. (2006, February
17). On making the right choice: The deliberation-without attention
effect. *Science*, *311*, 1005–1007.

- Duckworth, K.L., Bargh, J.A., Garcia, M., & Chaiken, S. (2002). The automatic evaluation of novel stimuli. *Psychological Science, 6*, 515-519.
- Dunning, D., Griffin, D. W., Milojkovic, J. H., & Ross, L. (1990). The overconfidence effect in social prediction. *Journal of Personality and Social Psychology, 58*, 568-592.
- Epstein, S. (1991). Cognitive-experiential self-theory: An integrative theory of personality. In R. Curtis (Ed.), *The self with others: Convergences in psychoanalytic, social, and personality psychology* (pp. 111–137).
- Evans, J. (2004). History of the dual process theory in reasoning. In K. I. Manktelow & M. C. Chung (Eds.), *Psychology of reasoning: Theoretical and historical perspectives* (pp. 241–266). Hove, England: Psychology Press.
- Evans, J. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology, 59*, 255-278.
- Fazio, R. H. (2001). On the automatic activation of associated evaluations: an overview. *Cognition & Emotion, 15*, 115-141.
- Fazio, R. H. (2007). Attitudes as object-evaluation associations of varying strength. *Social Cognition, 25*, 603-637.
- Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology, 54*, 297-327.

- Ferguson, M. J., & Fukukura, J. (2012). Like and dislikes: Attitudes from a social cognitive perspective. Invited chapter to appear in S. Fiske, & C. N. Macrae (Eds.), *Sage Handbook of Social Cognition*.
- Fiedler, K. (1988). Emotional mood, cognitive style, and behavior regulation. In K. Fiedler & J. P. Forgas (Eds.), *Affect, cognition and social behavior* (pp. 100-119). Gottingen, Germany: Hogrefe.
- Fitzsimons, G. M., & Bargh, J. A. (2004). Automatic Self-Regulation. In R. F. Baumeister, & K. D. Vohs (Eds.) *Handbook of Self-Regulation*. New York: Guilford Press.
- Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin*, *132*, 692-731.
- Gawronski, B., & Bodenhausen, G. V. (2011). The associative-propositional evaluation model: Theory, evidence, and open questions. *Advances in Experimental Social Psychology*, *44*, 59-127.
- Gawronski, B., & Sritharan, R. (2010). Formation, change, and contextualization of mental associations: Determinants and principles of variations in implicit measures. In B. Gawronski, & B. K. Payne (Eds.), *Handbook of implicit social cognition: Measurement, theory, and applications* (pp. 216-240). New York, NY: Guilford Press.

- Gawronski, B., LeBel, E. P., & Peters, K.P. (2007). What do implicit measures tell us? Scrutinizing the validity of three common assumptions. *Perspectives on Psychological Science*, 2, 181-193.
- Gigerenzer, G. (2007). *Gut feelings: The intelligence of the unconscious*. New York: Viking.
- Gigerenzer, G., & Regier, T. (1996). How do we tell an association from a rule? Comment on Sloman (1996). *Psychological Bulletin*, 119, 23-26.
- Gigerenzer, G., Todd, P. M., & the ABC Group. (1999). *Simple heuristics that make us smart*. New York, NY: Oxford University Press.
- Gilovich, T. (1991). *How we know what isn't so: The fallibility of human judgment in everyday life*. New York: Free Press.
- Gilovich, T., Griffin, D.W., & Kahneman, D. (2002). *The psychology of intuitive judgment: Heuristic and biases*. Cambridge: Cambridge University Press.
- Gladwell, M. (2005). *Blink: The power of thinking without thinking* (1st ed.). New York: Little, Brown and Co.
- Glockner, A., & Witteman, C. L. M. (2010). Beyond dual-process models: A categorisation of processes underlying intuitive judgement and decision making. *Thinking & Reasoning*, 16(1), 1-25.
- Gore, J., & Sadler-Smith, E. (2011). Unpacking Intuition: A Process and Outcome Framework. *Review of General Psychology*, 15(4), 304-316.

- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology, 74*, 1464-1480.
- Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the Implicit Association Test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology, 85*, 197-216.
- Gregg, A. P., Seibt, B., & Banaji, M. R. (2006). Easier Done Than Undone: Asymmetry in the malleability of implicit preferences. *Journal of Personality and Social Psychology, 90*(1), 1-20.
- Griffin, D. W., Dunning, D., & Ross, L. (1990). The role of construal processes in overconfident predictions about the self and others. *Journal of Personality and Social Psychology, 59*, 1128-1139.
- Hall, P. (1980). *Great planning disasters*. London, England: Weidenfeld & Nicolson.
- Hammond, K. R., Hamm, R. M., Grassia, J., & Pearson, T. (1987). Direct comparison of the efficacy of intuitive and analytical cognition in expert judgement. *IEEE Transactions on Systems, Man and Cybernetics, 17*, 753-767.
- Hassin, R. R., Bargh, J. A., & Zimmerman, S. (2009). Automatic and Flexible: The case of non-conscious goal pursuit. *Social Cognition, 27*(1), 20-36.
- Hermer-Vazquez, L., Hermer-Vazquez, R., Rybinnik, I., Greebel, G., Keller, R., Xu, S., Chapin, J.K. (2005). Rapid learning and flexible memory in "habit" tasks

- in rats trained with brain stimulation reward. *Physiology & Behavior*, 84(5), 753-9.
- Hertwig, R., Hoffrage, U., & Martingnon, L. (1999). Quick estimation: Letting the environment do the work. In G. Gigerenzer, P. M. Todd, & A. B. C. Research Group (Eds.), *Simple heuristics that make us smart* (pp. 37–58). New York: Oxford University Press.
- Hilliard, S., Nguyen, M., & Domjan, M. (1997). One-trial appetitive conditioning in the sexual behavior system. *Psychonomic Bulletin and Review*, 4, 237–241.
- Hofmann, W., Gschwendner, T., Nosek, B. A., & Schmitt, M. (2005). What moderates implicit-explicit consistency? *European Review of Social Psychology*, 16, 335-390.
- Hogarth, R. M. (2001). *Educating intuition*. Chicago: University of Chicago Press.
- Hogarth, R. M. (2010). Intuition: A challenge for psychological research on decision-making. *Psychological Inquiry*, 21(4), 338-353.
- Hogarth, R. M., Portell, M., & Cuxart, A. (2007). What risks do people perceive in everyday life? A perspective gained from the experience sampling method (ESM). *Risk Analysis*, 27, 1427– 1439.
- Hogarth, R. M., Portell, M., Cuxart, A., & Kolev, G. I. (2010). Emotion and reason in everyday risk perception. *Journal of Behavioral Decision Making*.

- Ijzerman, H., & Semin, G. R. (2009), The Thermometer of Social Relations: Mapping Social Proximity on Temperature. *Psychological Science*, 20(10), 1214–20.
- Inbar, Y. (2008). *Intuitions about intuitive insight and intuitive choice*. Thesis (Ph.D), Cornell University, August, 2008.
- Inbar, Y., Cone, J., & Gilovich, T. (2010). Intuitions about intuitive insight and intuitive choice. *Journal of Personality and Social Psychology*, 99, 232-247.
- Isen, A. M. (1987). Positive affect, cognitive processes, and social behavior. In B. Leonard (Ed.), *Advances in experimental social psychology*, Vol. 20. (pp. 203-253). San Diego, CA: Academic Press.
- Isen, A. M., Means, B., Patrick, R., & Nowicki, G. (1982). Some factors influencing decision-making strategy and risk-taking. In M. S. Clark & S. T. Fiske (Eds.), *Affect and cognition* (pp. 243-261). Hillsdale, NJ: Erlbaum.
- Jostmann, N. B., Lakens, D., & Schubert, T. W. (2009). Weight as an Embodiment of Importance. *Psychological Science*, 20 (9), 1169–74.
- Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, 58, 697-720.
- Kahneman, D. (2011). *Thinking fast and slow*. New York: Farrar, Strauss, Giroux.
- Kahneman, D., & Tversky, A. (1982). On the study of statistical intuitions. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 493–508). Cambridge, UK: Cambridge University

Press.

Keren, G., & Schul, Y. (2009) Two is not always better than one: A critical evaluation of two-system theories. *Perspectives on Psychological Science*, 4, 533-550.

Kirkpatrick, L. A., & Epstein, S. (1992). Cognitive-experiential self-theory and subjective probability: Further evidence for two conceptual systems. *Journal of Personality and Social Psychology*, 63, 534-544.

Klein, G. A. (1999). Sources of power: How people make decisions. Cambridge, MA: MIT Press.

Klein, G. A. (2003). *Intuition at work*. New York, NY: Currency and Doubleday.

Kruglanski, A. W., & Thompson, E. P. (1999). Persuasion by a single route: A view from the unimodel. *Psychological Inquiry*, 10, 83-109.

Krumhansl, C. L. (1997). An exploratory study of musical emotions and psychophysiology. *Canadian Journal of Experimental Psychology*, 51, 336-353.

Le Doux, J. E. (1996). *The emotional brain*. New York, NY: Simon and Schuster.

LeDoux, J.E. (2000). Emotion circuits in the brain. *Annual Review of Neuroscience*, 23, 155–184.

Mackie, D. M., & Worth, L. T. (1989). Cognitive deficits and the mediation of positive affect in persuasion. *Journal of Personality and Social Psychology*, 57, 27-40.

- McClelland, J.L., & Rumelhart, D.E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, *114*, 159–188.
- McConnell, A. R., Rydell, R. J., Strain, L. M., & Mackie, D. M. (2008). Forming implicit and explicit attitudes toward individuals: Social group association cues. *Journal of Personality and Social Psychology*, *94*, 792-807.
- Meehl, P.E. (1954). *Clinical versus statistical prediction: A theoretical analysis and a review of the evidence*. Minneapolis: University of Minnesota.
- Myers, D. G. (2002). *Intuition: Its powers and perils*. New Haven: Yale University Press.
- Nisbett, R., & Wilson, T. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, *84*, 231-259.
- Olson, M. A., & Fazio, R. H. (2001). Implicit attitude formation through classical conditioning. *Psychological Science*, *12*(5), 413-417.
- Oppenheimer, D.M., Meyvis, T., Davidenko, N. (2009). Instructional Manipulation Checks: Detecting Satisficing to Increase Statistical Power. *Journal of Experimental Social Psychology*, *45*, 867-872.
- Otten, S., & Wentura, D. (1999). About the impact of automaticity in the minimal group paradigm: Evidence from affective priming tasks. *European Journal of Social Psychology*, *29*, 1049-1071.

- Payne, B. K., Cheng, C. M., Govorun, O., & Stewart, B. (2005). An inkblot for attitudes: Affect misattribution as implicit measurement. *Journal of Personality and Social Psychology, 89*, 277-293.
- Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 19, pp. 123–205). New York, NY: Academic Press.
- Petty, R. E., & Fazio, R. H., & Briñol, P., (2009). The new implicit measures: An overview. In R. E. Petty, R. H. Fazio, & P. Briñol (Eds.), *Attitudes: Insights from the new implicit measures* (pp. 3-18). New York: Psychology Press.
- Petty, R. E., & Wegener, D. T. (1999). The elaboration likelihood model: Current status and controversies. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 37–72). New York, NY: Guilford Press.
- Petty, R. E., Tormala, Z. L., Briñol, P., & Jarvis, W.B.G. (2006). Implicit ambivalence from attitude change: An exploration of the PAST Model. *Journal of Personality and Social Psychology, 90*, 21-41.
- Plessner, H., Betsch, T., Schallies, E., & Schwierer, C. (2008). Automatic online formation of implicit attitudes towards politicians as a basis for intuitive voting behaviour. In H. Plessner, C. Betsch and T. Betsch (eds.), *Intuition in Judgment and Decision- Making* (pp. 107-117). Mahwah: Lawrence Erlbaum Associates.
- Poldrack, R.A., & Foerde, K. (2007). Category learning and the memory systems debate. *Neuroscience and Biobehavioral Reviews, 32*, 197–205.

- Pronin, E., & Wegner, D. M. (2006). Manic thinking: Independent effects of thought speed and thought content on mood. *Psychological Science, 17*, 807-813.
- Pronin, E., Jacobs, E., Wegner, D. M. (2008). Psychological effects of thought acceleration. *Emotion, 8*, 597-612.
- Reyna, V. F., & Brainerd, C. J. (2011). Dual processes in decision making and developmental neuroscience: A fuzzy-trace model. *Developmental Review, 31*, 180-206.
- Rozin, P., & Royzman, E.B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review, 5*, 296-320.
- Rutishauser, U., Mamelak, A. N., & Schuman E. M. (2006). Single-trial learning of novel stimuli by individual neurons of the human hippocampus-amygdala complex. *Neuron, 49*, 805-813.
- Rydell, R. J., & McConnell, A. R. (2006). Understanding implicit and explicit attitude change: A systems of reasoning analysis. *Journal of Personality and Social Psychology, 91*, 995-1008.
- Rydell, R. J., McConnell, A. R., Mackie, D. M., & Strain, L. M. (2006). Of two minds: Forming and changing valence-inconsistent implicit and explicit attitudes. *Psychological Science, 17*, 954–958.
- Rydell, R., McConnell, A. R., Strain, L. M., Claypool, H. M., & Hugenberg, K. (2007). Implicit and explicit attitudes respond differently to increasing

- amounts of counterattitudinal information. *European Journal of Social Psychology*, 37, 867-878.
- Schooler, J. W., & Engstler-Schooler, T. Y. (1990). Verbal overshadowing of visual memories: Some things are better left unsaid. *Cognitive Psychology*, 22, 36-71.
- Schwarz, N. (1990). Feelings as information: Information and motivational function of affective states. In E. T. Higgins & R. M. Sorrentino (Eds.), *Handbook of motivation and cognition* (Vol. 2, pp. 527-561). New York: Guilford.
- Schwarz, N., & Bless, H. (1991). Happy and mindless, but sad and smart? The impact of affective states on analytical reasoning. In J. P. Forgas (Ed.), *Emotion and social judgment* (pp. 55-72). New York: Pergamon Press.
- Schwarz, N., & Clore, G. L. (1983). Mood, misattribution, and judgments of wellbeing: Informative and directive functions of affective states. *Journal of Personality and Social Psychology*, 45, 513-523.
- Schwarz, N., Bless, H., & Bohner, G. (1991). Mood and persuasion: Affective states influence the processing of persuasive communications. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 24, pp. 161-199). Orlando, FL: Academic Press.
- Shah, A. K., & Oppenheimer, D. M. (2008). Heuristics made easy: An effort-reduction framework. *Psychological Bulletin*, 134(2), 207-222.

- Simon, H. A. (1955). A behavioural model of rational choice. *Quarterly Journal of Economics*, 69, 99–118.
- Simon, H. A. (1987). Making management decisions: The role of intuition and emotion. *Academy of Management Executive*, Feb., 57-64.
- Simon, H. A., & Chase, W. G. (1973). Skill in chess. *American Scientist*, 61, 393-403.
- Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3–22.
- Sloman, S.A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3–22.
- Smith, E. R., & Conrey, F. R. (2007) Agent-Based Modeling: A New Approach for Theory Building in Social Psychology. *Personality and Social Psychology Review*, 11(1), 87-104.
- Smith, E.R., & DeCoster, J. (2000). Dual process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review*, 4, 108–131.
- Squire, L.R., & Zola, S.M. (1996). Structure and function of declarative and nondeclarative memory systems. *Proceedings of the National Academy of Sciences*, 93, 13515–13522.
- Stanovich, K. E. (2004). *The Robot's Rebellion: Finding Meaning in the Age of Darwin*. Chicago: University Of Chicago Press.

- Stanovich, K. E. (2009). Distinguishing the reflective, algorithmic, and autonomous minds: Is it time for a tri-process theory? In J. Evans & K. Frankish (Eds.), *In two minds: Dual processes and beyond* (pp. 55-88). Oxford: Oxford University Press.
- Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences*, *23*, 645–665.
- Strack, F. & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review*, *8*(3), 220-247.
- Tesser, A. (1988). Towards a self-evaluation maintenance model of social behavior. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 21, pp. 181–227). San Diego: Academic Press.
- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgments. *Psychological Review*, *90*, 293-315.
- Vallone, R. P., Griffin, D. W., Lin, S., & Ross, L. (1990). Overconfident prediction of future actions and outcomes by self and others. *Journal of Personality and Social Psychology*, *58*(4), 582-592.
- Verosky, S. C., & Todorov, A. (2010). Generalization of affective learning about faces to perceptually similar faces. *Psychological Science*, *21*, 779-785.

- Weiss, H. M., & Cropanzano, R. (1996). Affective events theory: A theoretical discussion of the structure, causes and consequences of affective experiences at work. *Research in Organizational Behavior, 18*, 1–74.
- Wigley, S. (2007). Automaticity, Consciousness and Moral Responsibility. *Philosophical Psychology, 20*(2), 209–225.
- Williams, L. E., & Bargh, J. A. (2008), Experiencing Physical Warmth Promotes Interpersonal Warmth. *Science, 322* (5901), 606–7.
- Wilson, T. D. (2002). *Strangers to ourselves: Discovering the adaptive unconscious*. Cambridge, MA: Harvard University Press.
- 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.
- Wilson, T. D., Lindsey, S., & Schooler, T. (2000). A model of dual attitudes. *Psychological Review, 107*, 101-126.
- Worth, L. T., & Mackie, D. M. (1987). Cognitive mediation of positive affect in persuasion. *Social Cognition, 5*, 76-94.
- Yin, H. H. & Knowlton, B. J. (2006). The role of the basal ganglia in habit formation. *Nature Reviews Neuroscience, 7*, 464–476.
- Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist, 35*, 151–175.
- Zhong, C., & Liljenquist, K. (2006), Washing Away Your Sins: Threatened Morality and Physical Cleansing. *Science, 313* (5792), 1451–52.

