

The Goal of Cognitive Consistency: Priming and Applications to Judgment and Decision-
Making

Abstract

A fundamental criterion of judgment is consistency among beliefs. This cognitive consistency has a long history, but limited methods. Treating cognitive consistency as a goal, we present a new method for priming it. Three studies validate the method by using a word naming task, by showing an increase in predecisional information distortion, and by participants' direct reports of greater goal activation. The method is then used to verify the role of the consistency goal in three diverse judgment phenomena. Priming cognitive consistency increases the search for postdecisional supporting information (selective exposure to information), the agreement between preference and prediction (wishful thinking), and the adjustment of an implicit attitude assessed by the IAT to conform to the corresponding explicit attitude. Thus, the cause of these and possibly other judgment phenomena is not only a well-recognized desired outcome (motivated reasoning) but also the purely cognitive process of trying to make beliefs consistent.

Key-words: Goals, cognitive consistency, cognitive dissonance, wishful thinking, IAT, implicit attitudes, predecisional distortion of information, selective exposure

The Goal of Cognitive Consistency: Priming and Applications to Judgment and Decision-Making

Cognitive consistency is a fundamental topic with a rich history and intensive involvement in many judgment phenomena. It was the focus of substantial research in the 1960's, particularly by social psychologists studying cognitive dissonance (e.g., Abelson et al., 1968). However, interest in it appeared to decline so much that by 1983 Abelson could title an article, "Whatever became of consistency theory?" Nonetheless, research in cognitive consistency progressively reappeared under alternative labels like balance (Cvencek, Greenwald, & Meltzoff, 2012; Langer, Walther, Gawronski, & Blank, 2011), coherence (Holyoak & Simon, 1999; Glockner, Betsch, & Schindler, 2010; Thagard, 2000), compatibility (Goertzel, 1994), continuity (Zimbardo, 1999), and equilibrium (Cramer et al., 2012). The diversity of labels was compounded by a dispersal of interest across several subareas of behavioral science, such as belief systems (McGuire & McGuire, 1991), interpersonal relations (Walther & Weil, 2010), legal judgments (Engel & Glockner, 2012; Pennington & Hastie, 1988; Simon, Snow, & Read, 2004), reasoning (Johnson-Laird, Girotto, & Legrenzi, 2004; Thagard, 2006), and social cognition (Greenwald, Pickrell, & Farnham, 2002; Heine, Proulx, & Vohs, 2006; Read & Marcus-Newhall, 1993). The underlying role of cognitive consistency across these subareas is increasingly acknowledged, and with that recognition has come a return to greater research prominence (Gawronski & Strack, 2012).

Cognitive Consistency as a Goal

What is cognitive consistency? It is coherence among beliefs, including the case of the fit between an existing belief and new information that is essential to many judgment and decision tasks. However, this definition amounts to little more than a rephrasing of cognitive consistency. A deeper answer is hindered by the existence of multiple competing conceptualizations that have not been reconciled. For some cognitive consistency is a “*fundamental principle of information processing*” (Gawronski & Strack, 2012, p. 3) or, equivalently, a basic property of belief systems (McGuire, 1968). Others see it as a structure of relations among cognitive elements or as a measurable property of such a structure (Read & Simon, 2012). For a third group of researchers cognitive consistency is a “*means of disconfirming or refuting a hypothesis or a knowledge claim*” (Kruglanski and Shteynberg, 2012, p. 250). This notion of cognitive consistency as a means to an end is similar to seeing it as a mindset, that is, as a procedure learned and retained in memory that achieves a desired goal (Wyer and Xu, 2010). Finally, cognitive consistency has also been studied as a goal (Russo, Carlson, Meloy, & Yong, 2008), that is, a desired end-state (Fishbach and Ferguson, 2007). This last is the perspective that we adopt. In doing so, we in no way imply that the alternative views of cognitive consistency are invalid. However, reconciling the different conceptualizations of cognitive consistency is well beyond the present work.

Why Cognitive Consistency Now?

The current interest in cognitive consistency is prompted by two factors. First and recognized from the beginning, the fundamental nature of cognitive consistency suggests its potential role in influencing a broad range of phenomena. An individual’s ability to understand the external environment and, more importantly, to predict its opportunities and risks depends on sensing consistent relations among various cues and consequences. That is, the skill of detecting

consistencies is necessary to function in diverse environments. Consider a single topic like belief systems. Cognitive consistency might help to understand their structure (Read & Simon, 2012), the need to maintain a justification for a belief (Nickerson, 1998), and the reconciliation of new information with existing beliefs (Nickerson, 2008). As Nickerson (2008, p. 59) states, “*Surely everyone would agree that beliefs should be consistent with whatever evidence one has*”.

A second reason for exploring the role of cognitive consistency in various judgment and decision phenomena is the availability of newer methods for studying goals (Bargh et al. 2001; Ferguson & Porter, 2009; Förster, Liberman, & Friedman 2007). These methods can be used to test for the hypothesized influence of a particular goal on any phenomenon. In contrast, a case can be made that one reason for the declining activity in cognitive consistency after the 1960’s was the lack of general methods for manipulating the activation of consistency. Activating cognitive consistency in those early experiments often required manipulations that were task-specific. For instance, in the forced compliance paradigm, priming cognitive consistency typically required a participant to write a counterattitudinal statement, with the consistency-driven consequence being a change in that same attitude (Festinger & Carlsmith, 1959). Instead, the new methods activate goals independently of the subsequent focal task. Just as important to their validity, the new priming tactics can be implemented without the experimental participants’ awareness that a goal has been activated, thereby eliminating demand effects (Custers & Aarts 2010). Although the new methods have focused mainly on goal priming, some offer ways to assess the activation levels of goals (Carlson, 2001; Carlson, Tanner, Meloy, & Russo, 2012). Such assessments can both validate the success of the priming methods and measure goals’ activation levels across tasks and time. The availability of these methods enables tests of the goal

of consistency as a driver of judgment phenomena that could not have been performed in the past. Their use is essential to the current research.

The present work has two complementary objectives. First, we introduce and validate a method for nonconsciously priming cognitive consistency. Second, we apply that new method to several phenomena of human judgment, thereby illustrating the role of cognitive consistency in a range of tasks. Because we conceptualize the desire for cognitive consistency as a goal, our research necessarily builds on work on goals (e.g., Förster, Liberman, & Friedman, 2007; Baumgartner & Pieters, 2008).

A Method for Priming Cognitive Consistency as a Goal

The most familiar goal activation methods involve some form of semantic priming (Fishbach & Ferguson, 2007; Higgins, Rholes, & Jones, 1977; Neely, 1977). Participants are exposed to words related to a goal that, in turn, lead them to engage in goal-congruent behavior. For instance, when participants are asked to solve multiple scrambled sentences that include words related to hostility, they subsequently exhibit hostile behavior (Srull & Wyer, 1979). Semantic priming has already been successfully applied in research on consistency (Russo, Carlson, Meloy, & Yong, 2008; Experiment 1). However, we note that the power of semantic priming to actually drive a motivational state is disputed (Forster, Liberman, & Friedman, 2007; Nelissen, Dijk, & de Vries, 2005). It is possible that the observed impact of semantic priming on behavior occurs through an automatic perception-behavior link without needing an elevated motivational drive. That is, semantic methods may activate only the semantic network associated with a goal, including strongly associated actions. Partly for this reason we chose a different approach to priming the goal of cognitive consistency.

The priming of a motivational state requires a difference between the actual state and the desired state (Miller, Galanter, & Pribram, 1960; Sela & Shiv, 2009; Zeigarnik, 1967). Individuals who perform a task that creates such a discrepancy should experience an elevated level of a goal that can reduce that difference. Thus, we sought a task in which participants are faced with an inconsistency between two propositions that would be difficult to resolve. Further, both statements must be accepted as true so that their conflict cannot be resolved by denying one of them or dismissing a source. Then the expected lack of success in resolving the inconsistency should create the motivation to restore an acceptable level of consistency, that is, should activate the goal of cognitive consistency. Finally, this activation should carry over to an immediately subsequent focal task (Förster et al., 2007; Zeigarnik, 1967), thereby enabling a test of the role of the consistency goal in that target task.

Several attempts to find such a priming task led us to use a conundrum, a problem that requires participants to reconcile two seemingly inconsistent facts. One such conundrum is, “Why do we criticize traders for being overpaid but accept football players’ and movie stars’ salaries?” The inconsistency of beliefs highlighted by the conundrum and the expected inability to provide a fully satisfactory answer should activate the goal of restoring consistency within belief systems.

That activation level can be further increased by inserting a delay that frustrates the goal’s attainment (Bargh, Gollwitzer, Lee-Chai, Barndollar, & Troetschel, 2001; Förster et al., 2007). Thus, for any goal priming tactic, such as our conundrum, a delay of a few minutes caused by a filler task that prevents any progress toward goal achievement should produce an even higher level of activation¹. For this reason, our proposed priming method always consists of

¹ Although systematic demonstrations of the value of a delay in increasing a goal’s activation are relatively recent (Bargh et al., 2001; Bargh & Barndollar, 1996), it should be acknowledged that the benefit of a delay was

a conundrum to activate the goal of consistency and a filler task that delays goal attainment. This task-pair is then followed by one of several focal tasks that test a theoretical prediction about the role of cognitive consistency.

Note that our first objective is methodological. That is, we seek a method for reliably priming the goal of cognitive consistency, using whatever techniques are effective. We do not explore the basis of these techniques' success. Thus, we do not contrast the relative effectiveness of semantic and conundrum-driven priming or try to measure the benefit of a delay to increase goal activation. Our first objective is simply to find a goal priming tactic that works.

Validation of the Method: Studies 1, 2, 3, and 4

Studies 1-4 are designed to validate the conundrum-driven priming of the consistency goal by using three different tests for assessing the successful activation of a goal. The first is a semantic test: speeded naming of words that are related to consistency. If the goal of consistency is successfully activated by the conundrum (and the delay imposed by the filler task), then response times to the consistency-related words should be faster than when that goal is not activated (the control condition). The second study takes advantage of a known consequence of the increased activation of the consistency goal, the predecisional distortion of new information during a choice (Russo et al., 2008). If our priming successfully activates consistency, then we should be able to replicate past results and find that primed participants distort information more

discovered in the 1960's. Walster and Berscheid (1968), in the context of cognitive dissonance research, summarized the effects of time on cognitive consistency: "*A delay between the manipulation of dissonance and the measurement of the resulting dissonance reduction was incorporated into many dissonance studies. Researchers found that such a delay increased their predictive ability... in practice such a delay is often essential (p.599).*" We now recognize that this temporal pattern of an increase in activation as the result of a delay that prevents goal attainment is a standard characteristic of goals.

than control participants. Our third study both uses a different measure of goal activation, a direct report, and also examines more deeply the process at play in Study 2. Participants report the level of activation of the consistency goal using a method developed by Carlson et al. (2012). Successful priming should appear as higher reported activation of the consistency (but no other) goal. Further, higher reported levels of activation should mediate the impact of priming on information distortion. In addition, Study 3 also examines a potential effect of depletion. After these three different tests of the effectiveness of our conundrum-based priming method, Study 4 examines two potential confounds, negative affect and lowered self-esteem. Both are potential consequences of the inability to satisfactorily resolve the conundrum, and both might be seen to offer an alternative account of the effect of a greater activation of the consistency goal.

Study 1: Response Times

Study 1 uses an implicit assessment of goal activation based on the measurement of response times to naming visually presented words (Abad, Noguera, & Ortells, 2003; de Groot, 1985; Neely, 1991). If a goal is primed, then the network of concepts associated with the goal is activated, which results in faster reaction times to closely related words. Goal priming is verified by checking whether participants' relative speed in orally naming words related to the primed goal is increased when that focal goal is primed.

Participants

Twenty students took part in a computer-based experiment that was run in the laboratory of a large North American university. Each received a payment of \$5.

Method

The study design was a 2 (priming: consistency vs. control) x 2 (word category: consistency-related vs. neutral) mixed design. Priming was the between-subject factor; category of words was the within-subject factor.

Participants in the priming condition were given 3 minutes to solve the following conundrum: “Why do we criticize traders for being overpaid but accept football players’ and movie stars’ salaries?” The delay task was a 5-minute silent film, an extract of Charlie Chaplin’s “The Kid” (<http://www.youtube.com/watch?v=2MUBrClhgks>). Participants in the control condition did not have to work on the conundrum; they saw only the silent film.

After the priming phase, all participants were told that the next task would test their concentration skills. They were instructed to read aloud, as quickly as possible, each of a series of common words appearing on a computer screen. Response times were recorded with the expectation that, relative to the control condition, priming would reduce the time to name words related to consistency when compared to neutral words. The particular words comprising each category were drawn from Russo et al. (2008) who used a similar protocol. The consistency-related words were: agreement, coherence, compatible, congruence, consistent, and fitting. The neutral words, relatively unrelated to consistency, were: above, collection, deepen, kitchen, overcome, underline.

Participants began the naming task with five practice trials, during which no target words were used and reaction times were not recorded. They were then told that the actual task would begin. Each of the 12 test words was randomly presented and named, resulting in 240 total recorded response times (12 trials for each of 20 participants).

Results

Prior to statistical analysis, response times greater than 2000 ms and less than 150 ms were excluded from the analysis. These latencies typically reflect either a lack of attention or fatigue and are commonly removed from response time data sets (Bassili, 2003). This resulted in the loss of 5 of the 240 response times. For each participant, the remaining data were then averaged within category so that all participants had two mean response times, one for each category of words. These data were analyzed using a two-way mixed ANOVA, with word category as a within-subject factor and condition (control vs. consistency-priming) as a between-subject factor.

The test of the effectiveness of the priming method was a specific interaction that reflected faster RTs to consistency-related words (compared to neutral words) when the consistency goal was primed. A test of this interaction revealed it to be significant ($F(1, 18) = 8.80, p < .01$). The four mean RTs from the 2 x 2 design are depicted in Figure 1. For completeness, we note that the main effect of priming was not statistically reliable ($p > .05$), while that of word category was ($F(1, 18) = 19.63, p < .001$). Neutral words were named more quickly on average than were the consistency-related words. This latter difference was assumed to reflect word familiarity. Words such as “above” and “underline” were likely to be used more often than words such as “coherence” or “congruence”.

The presence of this interaction (Figure 1) demonstrated that our priming of cognitive consistency successfully increased the accessibility of consistency-related words, while having no significant effect on the accessibility of neutral words. The increased accessibility of consistency-related words supported the conclusion that cognitive consistency had been successfully activated.

Although such a test of priming is quite common, its interpretation is necessarily limited to the semantic domain. It is possible that only consistency-related words were activated and not the associated motivation to pursue cognitive consistency. Study 2 tests for the successful activation of the consistency goal by predicting an increase in a known effect of this goal.

Experiment 2: Information Distortion

As a second test of the ability of the conundrum-based method to activate the goal of cognitive consistency, we seek a common task whose behavior is known to be affected by this goal. Then activating the consistency goal should increase this behavior. To this end, we are fortunate that the consistency goal has been shown to drive an observable effect in binary choice (Russo et al., 2008). Thus, Experiment 2 uses the phenomenon known as the predecisional distortion of information as a behavioral measure of the activation of cognitive consistency. Information distortion is the biased evaluation of new information to support an emerging preference or other belief (Russo, Meloy, & Medvec, 1996). Specifically, during a choice process, as one alternative naturally emerges as the tentative leader, individuals typically interpret incoming information as favoring this leader more than they should. Because (semantic) priming of cognitive consistency has been shown to increase information distortion during a choice process (Russo et al., 2008), the efficacy of the new conundrum-based priming technique can be validated by finding a similar increase in a conundrum-primed group. Thus, there is a second test of the validity of the method that uses the known impact of consistency in the choice between two alternatives.

Participants

Participants were 70 students from a large North American university. They were paid \$5 for a 30-minute session in a laboratory and assigned randomly to either the priming or control condition.

Procedures

After the priming phase (which was identical to Study 1), participants chose between two options (resort hotels for a spring break vacation) based on five narrative attributes presented sequentially (e.g., Meloy, Russo, & Miller, 2006). After reading each attribute, participants were asked three questions. The first was a rating of the diagnosticity of the attribute on a 9-point scale, where 1 meant “strongly favors” one hotel, 9 meant “strongly favors” the other hotel, and 5 meant “favors neither hotel”. This response formed the basis of the measurement of information distortion. Second, participants were asked which alternative they would choose if they were to pick one at this stage, knowing that more information would be provided. This response identified the currently leading option. The third response was the participant’s confidence in this leading option.

Results

The calculation of information distortion required two steps. To begin, a participant’s estimate of an attribute’s diagnostic value (on the 9-point scale) was compared to an unbiased evaluation of the same information that was pre-tested to be approximately 5. This midpoint rating meant that the information favored neither hotel. Second, the absolute difference between those two values was signed positively if the attribute’s rating favored the leading alternative (which qualified as information distortion) and negatively if it favored the trailing one (which should only have occurred by chance). For example, the rating of the second attribute might have

been 6. If the alternative at the upper end of the scale was favored, then the absolute difference of 1 would be signed positively yielding an estimate of information distortion of +1. If the other hotel was favored, then the computed distortion would have been -1. Note that if the confidence in the leading hotel was zero, no reliable leader existed and no estimate of information distortion could be computed.

Because of the results of Russo et al. (2008), we expected participants with a higher activation level of the consistency goal to distort information more than those in the control condition. Confirming this expectation, the mean information distortion was significantly higher when consistency was activated ($M_{\text{priming}} = 1.15$, $SD_{\text{priming}} = 1.07$; $M_{\text{control}} = .40$, $SD_{\text{control}} = 1.49$; $t(68) = 2.41$; $p < .05$). Thus, we have a second, rather different, verification that our conundrum-based priming method can activate the goal of cognitive consistency.

We now turn to a third method, different yet again from the first two. We take advantage of new work on the direct assessment of goal activation. Further, we apply such an assessment technique to information distortion, the same phenomenon of Study 2. This should both replicate the effect of the consistency goal on information distortion and address consistency's mediation of that effect.

Study 3: Process via self-reports of goal activation

Previous research in the goals literature has shown that decision-makers, even if unaware of the source of goal priming, are still able to assess and report the goals that are active while making a decision (Carlson, 2001; Carlson et al., 2012). Such assessments can both validate the success of the priming methods and measure goals' activation levels across tasks and time.

The current availability of these methods enables tests of the goal of consistency as a driver of behavioral phenomena that could not have been performed in the past. Indeed, if participants are capable of reporting the level of the consistency goal during a task, then this self-reported activation should, in turn, explain any effect claimed to be due to priming the consistency goal. We investigate this process in Study 3.

The essence of Carlson's method (2001, 2012) for assessing the activation of goals in the act of decision-making is to interrupt participants during a task and ask them to record their level of activation of one or more goals. Note that the timing of this self-report is important. If participants are asked post-task, they must be able to remember something that was only partly available to conscious awareness at the time it occurred. Alternatively, participants might use the current (post-task) level of goal activation as a guide. However, when the task has been completed and the goal accomplished, its activation level has dropped substantially and might suggest that it had never been active at all (Förster, Liberman, & Friedman, 2007). Thus, it is essential to ask for the report of goal activation during the choice process, even when this requires interrupting that process. A second innovation of Carlson's method is to ask participants to report goal activation on a continuous scale. That is, the self-report method presumes that goal activation is not all-or-none, but lies on a continuum (Kouider, de Gardelle, Sackur, & Dupoux, 2010). Carlson et al. (2012) demonstrate that this sensitivity is essential to detecting differences in goal activation and that typical post-task reports of whether or not a goal was active may miss such differences entirely.

We apply this method of goal activation assessment in a choice task similar to that of Study 2. If our priming successfully activates cognitive consistency, participants should be able to report their use of a consistency-seeking strategy, which in turn should increase information

distortion.

In addition, this study also contributes to dismissing a potential effect of depletion on our results. In our first two studies, primed participants may have finished the first part of the experiment more tired than those in the control group. In turn, they may have reduced their effort in the subsequent criterion task (Baumeister, 2002; Johnson, 2007). Study 3 addresses this possibility in two ways. First, we match the effort exerted by both groups in the first part of the experiment, thereby eliminating differences between the primed and control conditions. Second, we directly measure and test the effect of our priming task on the reported use of an effort-saving strategy.

Participants

Through Amazon Mechanical Turk 200 native English speakers were recruited. Each received a payment of \$1. Participants were randomly assigned to the control condition or to the consistency-primed condition.

Procedure

Participants in the primed condition were given 3 minutes to solve the same conundrum as in Studies 1 and 2: “Why do we criticize traders for being overpaid but accept football players and movie stars' salaries?” They were then directed to a delay task of 3 minutes, in which they were instructed to answer a second open-ended question: “Describe one interesting thing you did last week. Your challenge is to enable the reader to experience, through your words, what you yourself did.” Participants in the control group were given 3 minutes to answer a question derived from the conundrum, in which no conflict had to be resolved, “Why do we criticize traders for being overpaid?” They were then directed to the same delay task as the primed group.

The objective of this new control procedure was to match the effort required from the participants in both groups in order to eliminate any potential impact of depletion.

The entire task was comprised of four segments: priming (or the question derived from the conundrum used in the control group), the filler task (to create a delay in the primed group but performed by all participants), informing participants of the four goals that they would have to report, and a consumer decision that was interrupted so that goal activation could be recorded. Following the priming phase and filler task, participants were given definitions of four goals: consistency, conserve effort, and two distracters appropriate to the choice task (memorize information, and separate the options). Carlson et al. (2012) argue that this pre-task familiarity with the goals is needed to ensure that all participants share the same meaning of each goal and can recognize them when they are interrupted and asked to report their activation levels. Each participant read the following instructions:

On the pages that follow, you will be asked to read and evaluate information that will help you make a decision between two products. During the choice process, you will be asked to report what strategies you are using to actually come to a decision.

To be sure that you fully understand the meaning of a question that will be asked later in the study, please take some time to read out loud the four following definitions:

- *Conserve Effort: Save effort when examining or making judgments about the options.*
- *Be consistent: See new information as consistent with information seen earlier.*
- *Separate the options: Exploit differences between the options to separate them from each other.*
- *Memorize: Commit the information seen to memory.*

Once you have read out loud the four definitions, you can start the study. Thanks in advance.

Participants then completed a choice task between two resort hotels, using only three of the five product attributes from Study 2. After both the second and third attributes, they reported the activation level of each of the four goals on a scale from 0 (*not at all*) to 100 (*as much as possible*). Two assessments were used to increase the sample size.

Results

The results are presented in three parts. First, were self-reported levels of the activation of the consistency goal higher in the primed condition? Second, did priming increase information distortion, thereby replicating Study 2's finding? Finally, could the first two results be linked by finding that the level of (self-reported) activation of the consistency goal mediated the magnitude of information distortion?

We found a reliable effect of priming on the mean self-reported level of the consistency goal in the expected direction ($M_{\text{consistency--conundrum}} = 66.58$, $SD_{\text{consistency--conundrum}} = 23.46$; $M_{\text{consistency--control}} = 58.57$, $SD_{\text{consistency--control}} = 25.69$; $t(1, 118) = 2.30$, $p < .05$). In clear contrast, priming the consistency goal had no effect on the reported activation level of the two distractors (separate the options and memorize; for both, $p > .05$). In addition, priming did not increase reliance on the "conserve effort" goal ($M_{\text{effort--conundrum}} = 50.29$, $SD_{\text{effort--conundrum}} = 26.49$; $M_{\text{effort--control}} = 48.20$, $SD_{\text{effort--control}} = 26.28$; $p < .05$). This latter null result, coupled with our roughly equating the time spent working by both primed and control groups, offered some evidence that people in the primed condition were not more depleted than participants in the control condition. This, in turn, suggested that depletion was unlikely to explain the results of our first two studies.

We then examined whether priming consistency would increase information distortion, thereby replicating the results of Study 2. Results confirmed a main effect of priming on the average distortion ($M_{\text{average ID-control}} = .83$, $SD_{\text{average ID-control}} = 1.99$; $M_{\text{average ID-conundrum}} = 1.31$, $SD_{\text{average ID-conundrum}} = 1.83$; $t(1, 118) = 2.51$, $p < .05$).

To determine whether self-reported activation of consistency mediated the relationship between condition and information distortion, we used a bootstrap method (Preacher & Hayes, 2008). Preacher and Hayes recommended the use of bootstrapping methods to test mediation to avoid limitations associated with other methods such as the Sobel test (Sobel, 1982) or the causal step approach (Baron & Kenny, 1986).

We recoded the independent variable (priming = 1 and control = -1), and specified it in the model as a categorical variable. Our dependent variable was the average distortion for each participant, and our potential mediating variable was the average self-report of the consistency goal activation. Results of this model are displayed in Figure 2. Without controlling for the mediator, priming had a significant positive effect on information distortion ($b = .24$, $SE = .10$, $p < .05$). Turning to the indirect path, priming had a significant positive effect on the self-report of consistency activation ($b = 4.01$, $SE = 1.73$, $p < .05$), and then self-reports of consistency activation had a significant positive effect on distortion ($b = .01$, $SE = .01$, $p < .05$). However, when controlling for the mediator, the direct impact of priming on distortion diminished and lost significance ($b = .18$, $SE = .14$, $p > .05$). We created bootstrap confidence intervals by using 10000 bootstrap samples. This yielded a 95% confidence interval of [.002; .140]. Because the interval excluded zero, it indicated that the self-reported use of a consistency strategy significantly mediated the effect of priming on information distortion.

Study 3 makes three contributions. First, it provides a third, quite distinct method, a direct self-report, that verifies the ability of the conundrum-based priming technique to activate the goal of cognitive consistency. Second, it strengthens the results obtained in Study 2 by demonstrating that differences in information distortion across participants are due to corresponding differences in the activation levels of the consistency goal. Third, it infirms the alternative explanation that depletion is a significant driver of our results.

The first three studies show, in their different ways, that the goal of cognitive consistency is activated by the conundrum-based priming method. However, we cannot claim that the priming method elicited no other effects that may have played some role in the observed results. Two such factors are negative affect and lowered self-esteem. Study 4 is designed to address both.

Study 4: Competing Explanations

The conundrum was designed to be difficult to resolve precisely so that it might activate the goal of seeking consistency. However, this same difficulty may have created other effects. Among these are negative affect and lowered self-esteem, both of which might be natural consequences of any frustration from not successfully resolving the conundrum. Study 4 is designed to test the effect of the conundrum-based priming method on both of these psychological states. We replicate Study 2 with the change that standard measures of both negative affect and self-esteem are added to the procedure. If the priming method activates only the goal of cognitive consistency and does not elicit negative affect or a decline in self-esteem, then there should be no difference between the priming and control conditions in these latter two.

However, because consistency should be activated, we should still replicate the effect of priming on information distortion.

Participants

One hundred and twenty native English speakers were recruited through Amazon Mechanical Turk. They were paid \$1 for their participation and randomly assigned to either the primed or control condition.

Method

Participants in the primed condition were given 3 minutes to solve the following conundrum: “Why do people buckle their seat belts more diligently in airplanes than in automobiles when statistically airplanes are safer than automobiles?” They were then directed to the delay task, in which they were instructed to read a text for 3 minutes. The text consisted of a paragraph extracted from Wikipedia, explaining where Patagonia is located and describing its geography. To match the time spent on these first tasks in both conditions, participants in the control group read the same text as did the primed group, but their version was extended to twice the length (about 6 minutes) with supplemental details on Patagonian geography.

After the priming phase, participants completed the two 10-item scales that compose the Positive Affect and Negative Affect Schedule (PANAS; Watson, Clark, & Tellegen, 1988). Participants also completed a 20-item scale measuring state self-esteem (SSES; Heatherton & Polivy, 1991). This scale was developed to be able to capture momentary fluctuations in self-esteem (e.g., “I feel inferior to others at this moment”). All items were scored on 7-point scales, from strongly disagree to strongly agree.

Following the completion of both the PANAS and SSES scales, participants chose between two options (running shoes) based on five narrative attributes presented sequentially.

Except for the switch of the focal product from resort hotels to running shoes, the procedure for assessing information distortion was identical to that in Study 2.

We expected to find the same effect of priming on information distortion. What was uncertain were whether that same priming also led to between-condition differences in negative affect and self-esteem.

Results

Results confirmed that the mean information distortion was significantly higher when consistency was activated ($M_{\text{priming}} = 1.53$, $SD_{\text{priming}} = 1.40$ vs. $M_{\text{control}} = .90$, $SD_{\text{control}} = 1.55$; $t(118) = 2.32$, $p < .05$).

We then turned to the analysis of the results of the SSEM and PANAS scales. There was no significant impact of priming on either self-esteem or negative affect. Participants in the primed condition had a slight tendency to show lower self-esteem, but this difference did not reach significance ($M_{\text{priming}} = 4.60$, $SD_{\text{priming}} = 1.07$ vs. $M_{\text{control}} = 4.87$, $SD_{\text{control}} = 1.07$; $t(118) = 1.38$, $p > .15$). Similarly, participants in the primed condition had a tendency to show more negative affect and less positive affect, as measured by the two 10-item PANAS scales, but neither of these effects reached significance ($M_{\text{NA-priming}} = 2.46$, $SD_{\text{NA-priming}} = 1.27$; $M_{\text{NA-control}} = 2.22$, $SD_{\text{NA-control}} = 1.09$; $M_{\text{PA-priming}} = 4.49$, $SD_{\text{PA-priming}} = .95$; $M_{\text{PA-control}} = 4.68$, $SD_{\text{PA-control}} = 1.07$; for both, $p > .15$). We also aggregated the affect data by subtracting the negative from the positive affect score to create a net affect score. The difference in this net affect score also failed to reach statistical significance ($M_{(\text{PA-NA})\text{-priming}} = 2.02$, $SD_{(\text{PA-NA})\text{-priming}} = 1.79$; $M_{(\text{PA-NA})\text{-control}} = 2.46$, $SD_{(\text{PA-NA})\text{-control}} = 1.86$; $t(118) = 1.29$, $p > .15$).

Conclusion

Study 4's main contribution is to eliminate two potential consequences of the conundrum method as explanations of any of the effects attributed to goal priming in Studies 1-3. There was no reliable effect of the priming manipulation on either negative affect or self-esteem. Thus, they could not have accounted for any of the three effects observed in those studies.

The first four studies yield the following conclusions. First, priming increases the semantic activation of concepts related to consistency. Second, priming increases information distortion, a phenomenon known to be caused by cognitive consistency. Third, priming increases information distortion through the activation of a consistency-seeking strategy in the act of decision-making. Fourth, priming can be detected by participants and directly reported as higher goal activation by using a sensitive method. Finally, the effects of priming on the criterion tasks cannot be explained by depletion, negative affect or lowered self-esteem. Based on these findings, we conclude that the conundrum-based priming method works. We can now turn to applying this method to judgment and decision phenomena that may benefit from a better understanding of their underlying process, particularly the role of the goal of cognitive consistency.

Applications of the Method: Studies 5, 6, and 7

From a large set of possible applications of the new method for priming the goal of cognitive consistency, we selected three. The first is the tendency to seek more confirming than disconfirming information after a decision, or selective exposure to information (Study 5). Selective exposure has been largely assumed in the past to be driven by cognitive consistency and, therefore, seemed like a straightforward initial application. The second phenomenon is the

effect of the desirability of an outcome or event on the judgment of its likelihood, or wishful thinking (Study 6). The prevalence of wishful thinking across contexts has been under debate in the last decade. Priming consistency may highlight one of the cognitive processes through which wishful thinking occurs, and thereby make a contribution to these debates. The more exploratory Study 7 extends the domain of application from explicitly stated beliefs to an implicit belief. This last study asks whether activation of the goal of consistency can increase the level of agreement between an explicit belief/attitude and one that is inaccessible to conscious reporting.

Study 5: Selective Exposure

One behavior that many would expect to be influenced by the desire for consistency is the post-decision reduction of cognitive dissonance (Frey, 1986). Biased information search to support a chosen option (Hart et al., 2009) is a special case of cognitive dissonance called selective exposure to information (Festinger, 1957). Because we use a choice task followed by a search for information that could support the chosen option, our study tests whether the desire for consistency is a driver of cognitive dissonance in the domain of information search. Although the desire for cognitive consistency has long been believed to underlie the reduction of cognitive dissonance, no one has directly verified its role by manipulating its level of activation.

Participants

One hundred and twenty native English speakers were recruited through Amazon Mechanical Turk and run online. Each participant was compensated 50 cents and randomly assigned to one of the two conditions, priming or control.

Procedures

Participants in the primed condition answered the following conundrum: “Why did international aid organizations receive from all around the globe \$500 million for relief of the Haiti earthquake but only \$5 million for the more devastating floods in Pakistan?” The delay task was the same as in Study 1, i.e., the same short movie. Participants in the control condition did not answer the conundrum.

Following the priming phase, the decision task required participants to make a choice between buying and not buying a digital camera based on an initial description and five user reviews. Participants could choose the five reviews from four sources: the company’s website, an online retailer website, a blog opposed to the camera, and a fan page. The company’s website and the fan page were described as likely to contain positive reviews. The blog and the online retailer were described as likely to contain negative reviews. Each review was written by an actual consumer, found online, and lightly edited for clarity and length as needed.

Our interest was only what participants did after they had made their buy/no-buy decisions, when they were instructed to select three more user reviews from the same set of four websites. The dependent variable was the proportion of these three reviews that was confirmatory. A review was coded “confirmatory” if the participant had decided to buy (not to buy) the camera and it was selected from one of the two more positive (negative) sources. The opposite held for “disconfirmatory”.

Results

The proportion of confirmatory reviews selected by the control group was .59. This value itself was significantly different from .50 ($t(59) = 2.66, p < .05$), replicating the standard selective exposure effect. More importantly, priming cognitive consistency significantly

increased the size of the selective exposure effect, from .59 to .70 ($SD_{\text{control}} = .27$ and $SD_{\text{priming}} = .29$; $t(118) = 2.16$; $p < .05$).

This last result demonstrated that the desire for cognitive consistency is at least one driver of the process that underlies the post-choice biased search for information. Fischer and Greitemeyer (2010) note that, even after 50 years of research on selective exposure, there are important open questions. We suggest that the role of consistency provides at least a partial answer to some of these questions.

Study 6: Wishful thinking

Study 6 examines the impact of cognitive consistency on wishful thinking. This is the influence of an outcome's desirability on the judgment of its expected likelihood. For a desired event, its likelihood is often elevated, while for undesired events it is depressed. Hastie (2001), in his review of 16 research problems to be solved in judgment and decision-making, calls for a better understanding of when and how wishful thinking occurs. He mentions that little evidence has been found for the reliable occurrence of this phenomenon across contexts. Similarly, recent research (e.g., Harris & Hahn, 2011, Krizan & Windschitl, 2007, Vosgerau, 2010, Windschitl, Smith, Rose, & Krizan, 2009) suggests that wishful thinking may not be as prevalent as initially thought. We suggest that cognitive consistency might be one the processes driving wishful thinking, and may in the future help to understand better under which circumstances this phenomenon occurs. The essence of the argument for why cognitive consistency might contribute to wishful thinking is that so long as there is an active desire to seek consistency

between two beliefs, there is a desire to reconcile the estimated likelihood of an event's outcome with the preference for that outcome.

Participants

The study was run online through Amazon Mechanical Turk, with 150 native English speakers. Two participants were disqualified for not completing the study. The remaining 148 participants were each paid 50 cents.

Procedures

Participants were first randomly assigned either to the control or primed conditions. Participants in the primed condition answered the following conundrum: "Why do medical volunteers in poor countries often have to use their own money to provide supplies while drug kingpins make millions and use that wealth to insulate themselves from arrest?" They were then directed to a delay task of 3 minutes, in which they were instructed to answer a second open-ended question: "Describe one interesting thing you did last week. Your challenge is to enable the reader to experience, through your words, what you yourself did." Participants in the control group only completed the second open-ended question.

After completing the priming task (or only the delay task for those in the control condition), all participants were asked to make a series of six predictions regarding the 83rd Academy Awards, held in February 2011. Note that the study was run in January 2011, near enough to the awards themselves that the actual list of nominees could be used.

The instructions were the following:

For each of the six major Academy Award categories listed below, please predict the winner to the best of your ability and knowledge. This is the nominee that you believe will win, regardless of whether they should win.

The six categories were: best actress, best actor, best supporting actor, best supporting actress, best director, and best picture. For instance, for the category best actor, the participants answered the following question:

Which nominee do you predict will win Best Actor? Colin Firth (The King's Speech); James Franco (127 Hours); Javier Bardem (Biutiful); Jeff Bridges (True Grit); Jesse Eisenberg (The Social Network).

After making all six predictions, participants were asked to indicate which nominees they would actually like to see win. The instructions, adapted for each category, were:

Now please tell us the nominee that you would most like to see win the Academy Award, regardless of their chance of winning. This is the nominee that you hope wins, that would most please you if they do win, completely independent of whether they actually will win.

Which nominee would you most prefer to win Best Actor? Colin Firth (The King's Speech); James Franco (127 Hours); Javier Bardem (Biutiful); Jeff Bridges (True Grit); Jesse Eisenberg (The Social Network).

Note that the likelihood judgments were requested before the preferences. This meant that the strength of the preferences was not increased by participant's being forced to consider them before they made their likelihood predictions. Instead, only the natural level of preferences should have influenced those judgments. We also believed, as is commonly assumed in the desirability bias, that preference judgments were more likely to be stable (i.e., independent of context) than likelihood judgments.

Results

The dependent measure was the number of matches between a participant's predictions and preferences, a number between 0 and 6. This amounted to an agreement score between

preferences and likelihoods. If the goal of cognitive consistency was a driver of wishful thinking, then priming consistency should have produced a higher agreement score.

The results confirmed this prediction. In the control condition the mean number of matches was 2.34 ($SD = 1.44$). In the priming condition, the corresponding mean increased to 3.05 ($SD = 1.58$). A mixed ANOVA was performed with type of prize category as the within-subject factor and priming as the between-subject factor. The main effect of priming was significant ($F(1, 147) = 8.30, p < .01$). In addition, a difference between prize categories was also observed ($F(5, 735) = 3.81, p < .005$), largely due to an overall higher agreement for the category Best Actress (dominated by Natalie Portman in *Black Swan* in 2011). However no interaction effect between prize category and priming condition was found ($F(5, 735) = 1.48, p > .15$). Thus, the effect of priming the consistency goal did not differ reliably across categories.

Just as for selective exposure in the previous study, wishful thinking was found to be driven, at least in part, by the goal of cognitive consistency. We do not claim that cognitive consistency is the entire explanation for wishful thinking, but that it is clearly a contributor in some situations. Further, if our assumption that consistency is a widespread goal is valid, then it may be a relatively common cause of this phenomenon.

Study 7: Implicit Beliefs

The previous tasks all involved consistency between beliefs that people are fully aware of. For instance, in wishful thinking, the two beliefs were a participant's preference among five alternatives and a prediction about those same five options. It seemed plausible, and our results

verified, that priming the goal of cognitive consistency would increase the agreement between such consciously held (and explicitly stated) beliefs.

We now ask whether cognitive consistency might operate on an implicit belief, one that individuals are much less consciously aware of. A common source of such beliefs is biases against a group of people, often racial or religious biases, that are widely considered to be morally unacceptable. Individuals may hold such biases implicitly but be reluctant to endure the social disapprobation of admitting to them publically or be reluctant to see themselves as “biased” persons by accepting them internally.

To test the possibility that our priming manipulation can affect the magnitude of an implicit bias, specifically to create “*a manifest change in concordance with self-reported attitude*” (Gregg, Seibt, & Greenwald, 2006, p. 15), we use the Implicit Association Test (“IAT”; Greenwald, McGhee, & Schwartz, 1998) and the bias against overweight individuals. Prior research has documented a consistent and pervasive anti-fat bias in settings as varied as education, employment, and healthcare (for a review, see Puhl & Brownell, 2001), as well as in the general population and in overweight and obese individuals themselves (see Watts and Cranney, 2009 for a review). Further, the anti-fat bias is observed in cases where explicit attitudes indicate little or no bias (e.g., Teachman & Brownell, 2001). Thus, we rely on a widespread presence of an implicit anti-fat bias and a widespread rejection of such a bias as socially unacceptable when it is required to be explicitly reported. The question is then whether priming cognitive consistency will alter the implicit anti-fat bias to bring it significantly closer to its explicit correspondent.

The notion that the consistency goal can shift the implicit anti-fat bias (towards a better alignment with the lower explicit bias) would not be universally accepted. One might reasonably

suggest that implicitly held beliefs are simply too inaccessible to be affected. In addition, implicit beliefs, especially those like an anti-fat bias, are considered to have been formed over time from multiple social experiences. As such, they are seen as relatively stable, not immutable but also not anywhere near as malleable as explicit beliefs (Gawronski & Bodenhausen, 2006). Further, there is empirical evidence that implicit and explicit beliefs are influenced along different pathways. Gawronski and Strack (2004) used the forced/induced compliance paradigm to test for an effect of cognitive dissonance reduction on both explicit and implicit attitudes. Each of their two studies found the usual effect of compliance on explicit attitudes, but no effect on the corresponding implicit attitudes (as assessed by the IAT). This is not to say that there are no empirical demonstrations of manipulated changes in implicit attitudes (e.g., Dasgupta & Greenwald, 2001). However, Gregg et al. (2006) found no empirical evidence of changes in implicit attitudes and have proffered a thoughtful account of the discrepancy between their more common negative result and the published examples of observed change.

Our test for the influence of the goal of cognitive consistency on an implicit belief is based on the agreement between the explicit and implicit forms of the anti-fat bias. We predict that when consistency is primed this agreement will increase. Such a test requires that participants report the explicit version of the anti-fat bias first, followed by the priming manipulation and then the assessment of the implicit bias. Given this experimental order of tasks, the predicted increase in agreement between the explicit and implicit version of the same attitude can be attributed only to a change in the latter. In contrast, if the implicit version were presented first (followed by priming and the explicit version), it is unlikely that the latter would change. The explicit anti-fat bias is likely to be more consciously, even strategically, driven by the goal

of social acceptance and, therefore, less susceptible to influence by the less dominant goal of cognitive consistency.

Participants

The experiment was run online through Amazon Mechanical Turk with native English speakers. Participants were compensated \$1 for their participation. Of the 175 participants who began the study, 16 (9%) failed to complete it due to technical difficulties (e.g., connectivity issues or use of an incompatible web browser or operating system). Of the remaining 159 participants, 6 participants with response latencies less than 300 ms on more than 10% of the IAT trials were excluded (Greenwald, Nosek, & Banaji, 2003)².

Procedure

Explicit Attitudes. All participants first completed an assessment of explicit attitudes toward overweight individuals, which was modified from Greenwald, Nosek, and Banaji (2003). This assessment was comprised of three components. The first measure was assessed by calculating the difference in reported feelings of warmth towards overweight and thin individuals (e.g., “*How warm or cold are your feelings towards fat/thin people?*”; 1=*Extremely warm* – 11=*Extremely cold*). The second measure was similarly obtained as the difference of self-reported identification with overweight and thin individuals (e.g., “*How much do you identify with people who are fat/thin?*”, 1=*Not at all* - 5=*Strongly*). A final measure assessed self-reported preference for overweight or thin individuals (1=*I strongly prefer thin people to fat people*; 7= *I strongly prefer fat people to thin people*). Responses to these three items were standardized and averaged to construct a single, combined measure of explicit anti-fat attitude/belief ($\alpha = .72$). Directionality was scored such that higher values indicated greater explicit bias.

² Reinstating these 6 participants did not affect the statistical significance of any result.

Although our test required that explicit attitude be assessed first, we note that Nosek, Greenwald, and Banaji (2005) found “little to no effect” on the magnitude of implicit and explicit attitudes as a function of presentation order. To confirm the necessity of presenting the explicit items first, we also ran a parallel study with 226 participants recruited and qualified in the same way³. However in this case, the explicit attitude was measured after the IAT had assessed implicit anti-fat bias. If our claim is correct, then there should be no effect of priming consistency on implicit attitude without prior completion of the explicit assessment.

Priming. Participants in the consistency prime condition responded to the following conundrum: “Why do people buckle their seat belts more diligently in airplanes than in automobiles when statistically airplanes are safer than automobiles?” Participants in the control condition read a neutral passage about the ancient Pharos lighthouse of Alexandria. The length of the passage (400 words) was designed to match the duration of the consistency prime (roughly 3 minutes). Participants in both conditions were then directed to a 3-minute extract from Charlie Chaplin's “The Gold Rush” (<http://www.youtube.com/watch?v=LZklHHg7QZI>) which served as the delay task.

IAT. Following the priming, participants completed the Fat-Thin IAT (copyright A. G. Greenwald, IAT Corp, 2001-2003, 2005) using Inquisit 3 Web Edition by Millisecond. The IAT consisted of seven blocks of sorting trials in which participants indicated the appropriate category of an image or word by pressing the “D” or “K” key of their computer keyboard (Greenwald, Nosek, & Banaji, 2003). In the first block (20 trials) participants classified silhouettes of body shapes as either “fat” or “thin”, followed by a second block (20 trials) in which participants classified various words (e.g. “love”, “peace”, “agony”, “evil”) as either

³ Of 227 participants who completed the study, 1 participant with response latencies less than 300 ms on more than 10% of IAT trials was eliminated.

“pleasant” or “unpleasant”. A third block of 20 trials, “compatible practice”, allowed participants to practice classifying images or words as either “thin or pleasant” (responses indicated by same key) or “fat or unpleasant” (same key response), followed by a fourth block of 40 trials of the same sorting conditions, the “compatible test”. In the fifth block, 40 fat /thin images were once again sorted, however this time using the reverse response key assignment⁴. In the sixth block of 20 trials (“incompatible practice”) participants practiced sorting images and words as either “fat or pleasant” or “thin or unpleasant”, eventually repeating this sorting combination in the seventh block (40 trials; “incompatible test”). The two pairs of “critical blocks”—blocks 3 and 6, and blocks 4 and 7—were used to compute participants’ IAT scores, as described below. Critical block orders were counterbalanced such that half of the participants first completed the compatible pairings in blocks 3 and 4 followed by incompatible combinations in blocks 6 and 7, while the remaining participants initially completed incompatible blocks followed by compatible blocks.

Note that the order of tasks was explicit assessment, priming (or reading the passage in the control condition), and the IAT. The priming manipulation was placed between the explicit and implicit attitude measurements so that it would not affect the former and would be as strong as possible in order to affect the latter.

Scoring. For each participant, the scoring procedure recommended by Greenwald et al. (2003) provided an individual “D-score,” a variant of Cohen’s *d*, which served as the measure of anti-fat bias from the IAT. In line with this procedure, trials with response latencies greater than 10,000 ms were eliminated. We subtracted participants’ average response latency during the compatible block (fat/unpleasant and thin/pleasant pairings) from the average response latency

⁴ As per Nosek et al. 2005, the number of trials in this block was increased from 20 to 40 to reduce pairing order effects.

during the incompatible block (fat/pleasant and thin/unpleasant pairings) for each of the two critical block pairs. Each of the resulting two difference values was then divided by the applicable pooled standard deviation (calculated independently for blocks 3 and 6, and for blocks 4 and 7), forming a separate quotient for each critical block pair. The D-score was computed as the average of these two quotients (Greenwald et al., 2003). D-scores were calculated such that more positive scores reflected greater implicit anti-fat bias.

Results

The test of the effect of priming consistency was based on an increase in the correlation between the explicit and implicit anti-fat attitudes. If consistency influenced the implicit attitude, then priming should have moved it closer to the explicit measure. The correlation between the two forms of the anti-fat attitude was .234 in the control condition, rising to .521 when the goal of cognitive consistency was primed. This increase was statistically significant based on the test for two independent samples, $z = 2.04$, $p = .041$. Both values were also reliably greater than zero, $p < .05$ for both.

We checked that it was necessary for the explicit attitude to be assessed before the implicit equivalent by switching the order of these measures in our parallel study ($n=226$). The correlation between the explicit and implicit anti-fat attitudes was .297 in the control condition and was .201 when consistency was primed. Their difference was not statistically significant based on the test for two independent samples ($z = -0.76$, $p = .447$) and was not even in the predicted direction.

Conclusion

These results show that the effect of cognitive consistency extends to beliefs that are less than fully accessible to conscious awareness, that is, to implicit beliefs. They add to the body of findings that such beliefs are not stable, but are malleable.

How might the activation of the consistency goal produce a shift in an implicit attitude such that it better conforms to one assessed via self-report? Adopting a dual-system model of information processing, Gawronski and Bodenhausen (2006) have argued that implicit and explicit attitude measures reflect two distinct processes. Specifically, while explicit measures primarily result from propositional processes concerned with the validity of information, implicit measures serve as a proxy for the activation of associations in memory, independent of the truth (or falsity) of those associations. Such a conceptualization is consistent with studies reporting a change in implicit attitudes in response to the formation (and activation) of new associations via an evaluative conditioning manipulation (Olson and Fazio 2006; Karpinski and Hilton 2001). In the absence of such a manipulation to create and activate new associations, priming a goal of cognitive consistency may have instead led to the preferential activation of *existing* associations that better conformed to the earlier provided self-report over those exhibiting less agreement. The greater relative activation of these consistent associations would then be reflected in an IAT score that tracks more closely the explicit measure. While this proposed mechanism is necessarily post hoc, its examination remains a promising direction for future research.

Our results may at first observation seem to run counter to Gawronski and Strack's (2004) failure of the forced/instructed compliance manipulation to affect an implicit attitude. These authors attributed their finding to the propositional elements characterizing both the cause of cognitive dissonance as well as the process of its reduction. In contrast, an effect on implicit attitudes as observed in the current research would involve, according to the dual-process model

described above, the sufficient activation of relevant associations in memory. While our findings may be driven by the activation of extensive existing associations consistent with a self-reported attitude, it may simply be that the induced compliance manipulation used by Gawronski and Strack (writing a counter-attitudinal essay) failed to exert any influence on implicit attitudes due to the relatively scarcer network of these related associations in memory.

General Discussion

Summary of results

Across seven studies we have introduced a method that primes cognitive consistency as a goal. Our first four studies demonstrate the efficacy of our priming method. Study 1 shows that the accessibility of consistency-related words is significantly increased when participants are exposed to a conundrum. Study 2 replicates a finding highlighted in previous literature: priming cognitive consistency increases information distortion, i.e., modifies how decision-makers evaluate information based on their current preferences. Study 3 provides a direct report of increased activation of the consistency goal due to priming. It also illuminates the process at work in Study 2, namely the mediating role of the consistency goal in driving information distortion. Further, it contributes to eliminating depletion as an explanation of the observed results. Study 4 then infirms two competing accounts of the priming manipulation, increased negative affect and reduced self-esteem.

We then applied the conundrum-based method to three diverse judgment phenomena: selective exposure to information, wishful thinking, and the agreement between an implicit and explicit attitude. Selective exposure was chosen based on the previous literature on cognitive

dissonance. Wishful thinking was chosen as a phenomenon whose ubiquity is under debate, and to which priming the consistency goal might make an empirical contribution. Study 7, using the IAT, verified that our conundrum-based priming method could influence a less conscious, implicit belief. Across these three applications, participants exposed to a conundrum, in comparison to control participants: selected more confirmatory information to fit their prior beliefs, matched more closely the estimated likelihood of an event's outcome with their preference for that outcome, and showed more congruence between their explicit and implicit beliefs.

Natural primings of cognitive consistency

What the ability of an experimental manipulation to prime cognitive consistency raises the issue of naturally occurring primes. One such naturally occurring prime may be the presence of a substantial conflict in a binary choice, such as when one attribute strongly favors alternative A while another just as strongly favors Alternative B. This situation may activate the goal of consistency and influence the rest of the decision process, for instance increasing information distortion as demonstrated in Experiment 3. As a second candidate for a naturally occurring prime, consider managers who are exposed to inconsistent information or information that conflicts with a current belief (e.g., market research that contradicts an expected consumer response). Our results suggest that, if such inconsistencies and conflicts activate the goal of cognitive consistency, there may be a nontrivial impact on subsequent information processing during judgments and decisions. These examples remain speculative until subjected to empirical tests.

Process versus Outcome Goals: Cognition versus Motivated Reasoning

Cognitive consistency is a process goal, one of those that are confined to the process itself and are satisfied, or frustrated, during that process (Carlson et al., 2008; Van Osselaer et al., 2005). Other examples of process goals are conserving effort (Payne, Bettman, & Johnson, 1990) and avoiding negative emotions (Luce, Payne, & Bettman, 2001). Process goals are distinct from goals linked to an outcome. For instance, in a decision, an outcome goal might be finding an automobile with high gas mileage or a food that is low in calories. Because outcome goals are task-specific, they tend to be numerous and varied. They also tend to be consciously reportable, as people are usually fully aware of pursuing such goals. For all of these reasons, substantial empirical attention has been devoted to outcome goals. In contrast, process goals are less numerous and often lie more in the background. If decision makers provided a concurrent verbal protocol, researchers would likely observe far more reports of outcome goals than of process goals. However, although not as numerous or varied, process goals often play a common role across different tasks. People strive to conserve effort and to minimize frustration whether the decision involves automobiles, foods or almost anything else. This generality of process goals suggests their value as objects of study

There is a second reason for making the process-outcome distinction. Phenomena, like selective exposure and wishful thinking, which have been classified as driven by outcome goals (and, therefore, included under motivated reasoning) may, in fact, be purely cognitive in nature. That is, although these phenomena are usually attributed to an outcome goal, such as wanting to support the chosen alternative (selective exposure) or to experience a more desired outcome (wishful thinking), the desire for cognitive consistency offers an alternative account that is purely cognitive. This speaks to the importance of cognitive consistency as broadly fundamental. It also offers a competing theoretical perspective on many judgment phenomena.

More phenomena

What other phenomena might be explained, at least in part, by the desire for cognitive consistency? We have presented cognitive consistency as a force for reconciling different beliefs. What activating consistency does is strengthen this force. Thus, one natural class of phenomena that might be driven by the desire for consistency is those where an existing belief confronts new information. Possibilities include the correspondence bias (Gilbert & Malone, 1995), the endowment effect (Kahneman, Knetsch, & Thaler, 1990), and the prior belief effect (Lord, Ross, & Lepper, 1979). Our hope is that what is necessarily speculation at present might, through the application of our priming method, be informed empirically.

Figure 1. Interaction between word category and priming

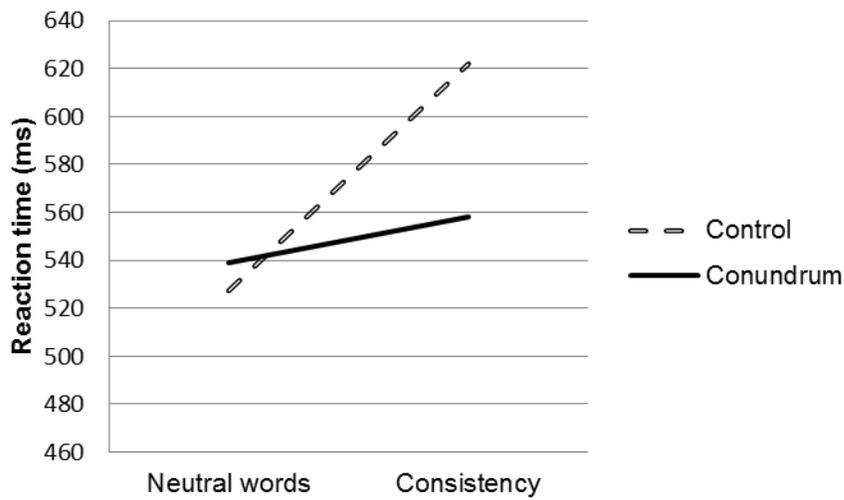
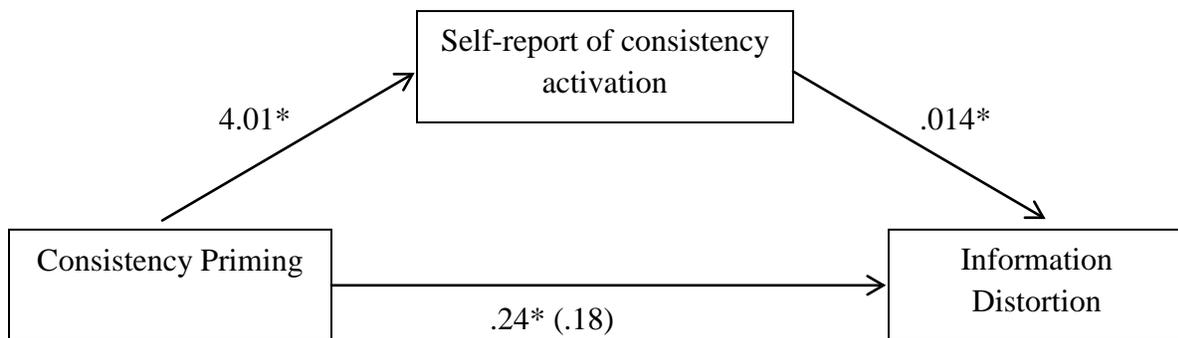


Figure 2. Mediation results for effect of condition on information distortion. The value in parentheses represents the direct effect of condition on information distortion, controlling for indirect effects. * $p < .05$



References

- Abad, M. J. F., Noguera C., & Ortells J.J. (2003). Influence of prime-target relationship on semantic priming effects from words in a lexical decision task. *Acta Psychologica, 113*, 283-295.
- Abelson, R.P. (1983). Whatever became of consistency theory? *Personality and Social Psychology Bulletin, 9*, 37-64.
- Abelson, R.P., Aronson, E., McGuire, W.J., Newcomb, E. T., Rosenberg, M.J., & Tannenbaum, R.H. (1968). *Theories of Cognitive Consistency: A Sourcebook*, Chicago, IL: Rand McNally.
- Bargh, J. A., & Barndollar, K. (1996). Automaticity in action: The unconscious as repository of chronic goals and motives. In P. M. Gollwitzer, & J. A. Bargh (Eds.), *The psychology of action: Linking cognition and motivation to behavior* (pp. 457-481). New York, NY: Guilford Press.
- Bargh, J. A., Gollwitzer, P. M., Lee-Chai, A., Barndollar, K. & Troetschel, R. (2001). The automated will: Nonconscious activation and pursuit of behavioral goal. *Journal of Personality and Social Psychology, 81*, 1014-1027.
- Baron, R. M., & Kenny, D. A. (1986), The moderator-mediator variable distinction in social psychological research: Conceptual, strategic and statistical considerations, *Journal of Personality and Social Psychology, 51*, 1173-1182.

- Bassili, J. N. (2003). Cognitive indices of social information processing. In A. Tesser & N. Schwartz (Eds.). *The Blackwell Handbook of Social Psychology: Intraindividual Processes* (pp. 68-88). Oxford, UK: Blackwell Publishing.
- Baumeister, Roy F. (2002), "Yielding to Temptation: Self-control Failure, Impulsive Purchasing, and Consumer Behavior," *Journal of Consumer Research*, 28, 670-676.
- Baumgartner, H., & Pieters, R. (2008). Goal-directed consumer behavior. In Haugtvedt, C. P., Herr, P. M., & Kardes, F. R. (Eds.), *Handbook of Consumer Psychology* (pp. 367-392), New York: Psychology Press.
- Carlson, K. A. (2001). In search of consumer goals. *Dissertation Abstracts International: Section B: The Sciences and Engineering, Vol. 62(7-B)*, Feb 2002, pp. 3368.
- Carlson, K. A., Janiszewski, C., Keeney, R. L., Krantz, D. H., Kunreuther, H. C., Luce, M. F., van Osselaer, S. M. J., & von Winterfeldt, D. (2008). A theoretical framework of goal-based choice and for prescriptive analysis. *Marketing Letters*, 19, 241-254.
- Carlson, K. A., Tanner, R. J., Meloy, M. G., & Russo, J. E. (2012). Catching goals in the act of decision making. Working Paper.
- Cramer, A. O. J., van der Sluis, S., Noordhof, A., Wicherts, M., Geschwind, N., Aggen, S. H., Kendler, K. S., & Borsboom, D. (2012). Dimensions of normal personality as networks in search of equilibrium: You can't like parties if you don't like people. *European Journal of Personality*, 26, 414-431.
- Custers, R. & Aarts, H. (2010). The unconscious will: How the pursuit of goals operates outside of conscious awareness. *Science*, 329, pp. 47-50.
- Cvencek, D., Greenwald, A. G., & Meltzoff, A. N. (2012). Balanced identity theory: Evidence for implicit consistency in social cognition. In Gawronski, B., & Strack, F. (Eds.),

- Cognitive consistency: A unifying concept in social psychology* (pp. 157–177). New York: Guilford Press.
- 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.
- De Groot, A. M. (1985). Word-context effects in word naming and lexical decision. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, *37A*, 281-297.
- Engel, C. & Glockner, A. (2012). Role-induced bias in court: An experimental analysis. *Journal of Behavioral Decision Making*, forthcoming.
- Ferguson, M. J., & Porter S. C. (2009). Goals and (implicit) attitudes. In G. B. Moskowitz and H. Grant (Eds.), *The Psychology of Goals* (pp. 454-460), New York: Guilford.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Evanston, IL: Row, Peterson.
- Festinger, L., & Carlsmith, J.M. (1959). Cognitive consequences of forced compliance. *Journal of Abnormal and Social Psychology*, *58*, 203–210.
- Fishbach, A., & Ferguson, M. J. (2007). The goal construct in social psychology. In A. W. Kruglanski and E. T. Higgins (Eds.), *Social Psychology: Handbook of basic principles* (pp. 490-515), New York: Guilford.
- Förster, J., Liberman, N., & Friedman, R. (2007). Seven principles of goal activation: A systematic approach to distinguishing goal priming from priming of non-goal constructs. *Personality and Social Psychology Review*, *11*, 211–233.
- Frey, Dieter (1986). Recent research on selective exposure to information. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (pp. 40-80), New York: Academic Press.

Gawronski, B., & Strack, F. (2004). On the propositional nature of cognitive consistency:

Dissonance changes explicit, but not implicit attitudes. *Journal of Experimental Social Psychology, 40*, 535-542.

Gawronski, B. and Strack, F. (2012). *Cognitive Consistency: A Fundamental Principle in Social Cognition*. New York, NY: The Guilford Press.

Gawronski, B. and Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin, 132*, 692-731.

Gilbert, D. T., & Malone P. S. (1995). The Correspondence bias. *Psychological Bulletin, 117*, 21-38.

Glockner, A., Betsch, T. & Schindler, N. (2010). Coherence shifts in probabilistic inference tasks. *Journal of Behavioral Decision Making, 23*, 439-462.

Goertzel, Ben (1994). *Chaotic Logic: Language, Thought, and Reality from the Perspective of Complex Systems Science*. New York: Plenum.

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*, 1-20.

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., Pickrell, J. E., & Farnham, S. D. (2002). Implicit partisanship: Taking sides for no reason. *Journal of Personality and Social Psychology, 83*, 367-379.

- 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.
- Harris, A. J. L., & Hahn, U. (2011), Unrealistic Optimism About Future Life Events: A Cautionary Note, *Psychological Review, 118*, 135-154.
- Hart, W., Albarracin, D., Eagly, A., Brechan, I., Lindberg, M. J., & Merrill, L. (2009). Feeling validated versus being correct: A meta-analysis of selective exposure to Information. *Psychological Bulletin, 135*, 555-588.
- Hastie R. (2001). Problems for judgment and decision-making. *Annual review of Psychology, 52*, 653-683.
- Heatherton T. F. & Polivy J. (1991). Development and validation of a scale for measuring self-esteem. *Journal of Personality and Social Psychology, 60*, 895-910.
- Heine, S. J., Proulx, T., & Vohs, K. D. (2006). The meaning maintenance model: On the coherence of human motivations. *Personality and Social Psychology Review, 10*, 88-110.
- Higgins, E. T., Rholes, W. S., & Jones, C. R. (1977). Category accessibility and impression formation. *Journal of Experimental Social Psychology, 13*, 141-154.
- Holyoak, K.J. and D. Simon (1999). Bidirectional reasoning in decision making by constraint satisfaction. *Journal of Experimental Psychology: General, 128*, 3-31.
- Johnson, E. J. (2002). Man, my brain is tired: Linking depletion and cognitive effort in choice. *Journal of Consumer Psychology, 18*, 14-16.
- Johnson-Laird, P.N., Girotto V., & Legrenzi P. (2004). Reasoning from inconsistency to consistency. *Psychological Review, 111*, 640-661.

- Kahneman D., Knetsch J. L., & Thaler R. H. (1990). Experimental tests of the endowment effect and the Coase theorem. *Journal of Political Economy*, *98*, 1325-1348.
- Karpinski, A., & Hilton, J.L. (2001). Attitudes and the Implicit Association Test. *Journal of Personality and Social Psychology*, *81*, 774–788.
- Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences*, *14*, 301-307.
- Krizan, Z., & Windschitl, P. D. (2007). The influence of outcome desirability on optimism. *Psychological Bulletin*, *133*, 95-121.
- Kruglanski, A. W., & Shteynberg, G. (2012). Cognitive consistency as a means to an end: How subjective logic affords knowledge. In B. Gawronski & F. Strack (Eds.), *Cognitive consistency: A Fundamental Principle in Social Cognition*. (pp. 245-264). New York: Guilford Press.
- Langer, T., Walther, E., Gawronski, B., & Blank, H. (2009). When linking is stronger than thinking: Associative transfer of valence disrupts the emergence of cognitive balance after attitude change. *Journal of Experimental Social Psychology*, *45*, 1232-1237.
- Lord, C. G., Ross L., & Lepper M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, *37*, 2098-2109.
- Luce, M. F., Payne, J. W., & Bettman, J. R. (2001). The emotional nature of trade-offs. In S. J. Hoch & H.C. Kunreuther (Eds.), *Wharton on Making Decision* (pp. 17-35), New York: Wiley.

- McGuire, W.J. (1968). Theory of the structure of human thought. In R.P. Abelson, E. Aronson, W.J. McGuire, T.M. Newcomb, M.J. Rosenberg, and R.H. Tannenbaum (Eds.). *Theories of Cognitive Consistency: A Sourcebook* (pp. 140-162), Chicago, IL: Rand McNally.
- McGuire, W.J., & McGuire, C.V. (1991). The content, structure, and operation of thought systems. In R. S. Wyer & T. K. Srull (Eds.), *Advances in Social Cognition* (Vol. IV, pp. 1-78), Hillsdale NJ: Erlbaum.
- Meloy, M. G., Russo, J. E., & Miller, E. G. (2006). Monetary Incentives and Mood. *Journal of Marketing Research*, 43, 267-275.
- Miller, G. A., Galanter, E. & Pribram, K. H (1960). *Plans and the Structure of Behavior*. New York: Holt, Rinehart & Winston.
- Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General*, 106, 226-254.
- Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264–336). Hillsdale, NJ: Erlbaum.
- Nelissen R. M. A., Dijkster A. J. M, & De Vries N. K. (2005). Limitations of semantic procedures for automatic goal activation. *Psychological reports*, 97, 675-689
- Nickerson, R. S. (1998). Confirmation bias: a ubiquitous phenomenon in many guises. *Review of General Psychology*, 2, 175–220.
- Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2005). Understanding and using the Implicit Association Test: II. Method variables and construct validity. *Personality and Social Psychology Bulletin*, 31, 166–180.

- Olson, M.A., & Fazio, R.H. (2006). Reducing automatically activated racial prejudice through implicit evaluative conditioning. *Personality and Social Psychology Bulletin*, 32, 421–433.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1990). The adaptive decision maker: Effort and accuracy in choice. In R. M. Hogarth (Ed.), *Insights in decision making: A tribute to Hillel J. Einhorn* (pp. 129-152), Chicago IL: The University of Chicago Press.
- Pennington, N., & Hastie, R. (1988). Explanation-based decision making: Effects of memory structure on judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 521- 533.
- Preacher K. J. & Hayes A. F. (2008), Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models, *Behavior Research Methods*, 40, 879-891.
- Puhl, R., & Brownell, K. D. (2001). Obesity, bias, and discrimination. *Obesity Research*, 8, 788–805.
- Read, S. J., & Simon, D. (2012). Parallel Constraint Satisfaction as a Mechanism for Cognitive Consistency. In B. Gawronski & F. Strack (Eds.), *Cognitive consistency: A Fundamental Principle in Social Cognition*. (pp. 66-86). New York: Guilford Press.
- Read, S.J. & Marcus-Newhall A. (1993). Explanatory coherence in social explanations: A parallel distributed processing account. *Journal of Personality and Social Psychology*, 65, 429-447.

- Russo, J. E., Carlson K. A., Meloy M. G., & Yong K. (2008). The goal of consistency as a cause of information distortion (2008). *Journal of Experimental Psychology: General*, *137*, 456-470.
- Russo, J. E., Meloy M. G., & Medvec V. H. (1996). The distortion of information during decisions. *Organizational Behavior and Human Decision Processes*, *66*, 102-111.
- Sela, A., & Shiv, B. (2009). Unraveling priming: When does the same prime activate a goal versus a trait? *Journal of Consumer Research*, *36*, 418-433.
- Simon, D., Snow, C.J., and Read, S.J. (2004). The redux of cognitive consistency theories: evidence judgments by constraint satisfaction. *Journal of personality and Social Psychology*, *86*, 814-837.
- Sobel, M. E., (1982). Asymptotic confidence intervals for indirect effects in structural equation models. In S. Leinhardt (Ed.), *Sociological Methodology* (pp. 290-312). Washington, DC: American Sociological Association.
- Srull, T. K., & Wyer, R. S. (1979). The role of category accessibility in the interpretation of information about persons: Some determinants and implications. *Journal of Personality and Social Psychology*, *37*, 1660-1672.
- Teachman B. A. & Brownell K. D. (2001). Implicit anti-fat bias among health professionals: Is anyone immune? *International Journal of Obesity*, *25*, 1525-1531.
- Thagard, P. (2000). *Coherence in thought and action*. Cambridge, MA: MIT Press
- Thagard, P. (2006). Evaluating explanations in law, science and everyday life. *Current Directions in Psychological Science*, *15* (3), 141- 145.

- Van Osselaer, S. M. J., Ramanathan, S., Campbell, M. C., Cohen, J. B., Dale, J. K., Herr, P. M., Janiszewski, C., Kruglanski, A. W., Lee, A. Y., Read, S. J., Russo, J. E., & Tavassoli, N. T. (2005). Choice based on goals. *Marketing Letters*, *16*, 335-346.
- Vosgerau, J. (2010). How prevalent is wishful thinking? Misattribution of arousal causes optimism and pessimism in subjective probabilities. *Journal of Experimental Psychology*, *139*, 32-48.
- Walster, E. & Berscheid, E. (1968). The effects of time on cognitive consistency. In R. P. Abelson, E. Aronson, W. J. McGuire, T. M. Newcomb, M. J. Rosenberg and P. H. Tannenbaum (Eds.), *Theories of cognitive consistency: a sourcebook*, Chicago: Rand McNally.
- Walther, E., and R. Weil (2010), Balance formation through retroactive construal. In Gawronski, B., & Strack, F. (Eds.), *Cognitive consistency: A unifying concept in social psychology*. New York: Guilford Press.
- Watson D., Clark L. A. and Tellegen A. (1988). Development and Validation of Brief Measures of Positive and Negative Affect: The PANAS Scales. *Journal of Personality and Social Psychology*, *54*, 1063-1070.
- Watts, K., and J. Cranney (1999). The nature and implications of implicit weight bias. *Current Psychiatry Reviews*, *5*; 2, 110-126.
- Windschitl P. D., Smith A. R., Rose J. P., and Krizan Z. (2010). The desirability bias in predictions: Going optimistic without leaving realism. *Organization Behavior and Human Decision Processes*, *111*, 33-47.

Wyer, R. S. and Xu A. J. (2010). The Role of behavioral mind-sets in goal-directed activity:

Conceptual underpinnings and empirical evidence. *Journal of Consumer Psychology*, 20, 107-112.

Zeigarnik, B.V. (1967). On finished and unfinished tasks. In W. D. Ellis (Ed.), *A sourcebook of*

Gestalt psychology, New York: Humanities Press

Zimbardo, P. G. (1999). Discontinuity Theory: Cognitive and Social Searches for Rationality and

Normality May Lead to Madness. *Advances in Experimental Social Psychology*, 31, 345-413.