Self-Serving Bias and Self-Deception

William von Hippel
University of New South Wales

Richard J. Shakarchi
Ohio State University

Jessica L. Lakin
Drew University

Abstract

Three experiments tested the hypothesis self-serving biases are self-deceptive in nature. Consistent with this hypothesis, Experiment 1 revealed that people who rated a task at which they succeeded as more important than a task at which they failed also cheated on a series of math problems, but only when they could excuse their cheating as unintentional. Experiment 2 replicated this finding and demonstrated that a self-report measure of self-deception did not predict this self-deceptive cheating. Experiment 3 replicated Experiments 1 and 2 and ruled out several alternative explanations. These experiments suggest that self-serving biases have a self-deceptive component, and that individual differences in self-deception can be measured.

KEY WORDS: Self-deception, self-serving bias, individual differences, cheating.

Correspondence should be addressed to William von Hippel, School of Psychology, University of New South Wales, Sydney 2052, Australia (w.vonhippel@unsw.edu.au).

Self-Serving Bias and Self-Deception

Most people associate themselves with desirable events and outcomes and distance themselves from undesirable events and outcomes (Baumeister, 1998; Schlenker, 1980; Sedikides, 1993). The most well known example of this self-serving bias (SSB) is the tendency to attribute success to internal factors such as ability and effort, and failure to external factors such as luck and task difficulty (Miller & Ross, 1975; Weary-Bradley, 1978; Zuckerman, 1979). Although there is an extensive literature on manifestations of the SSB and the conditions under which it emerges, the question of whether the SSB reflects self-presentation or self-deception remains open. There is ample evidence that self-serving biases are at least somewhat self-presentational (see Schlenker & Pontari, 2000), but it is more difficult to demonstrate that they are self-deceptive as well.

There are three classes of evidence that support the role of self-deception in self-serving biases. First, self-serving biases are shown even in the absence of an audience (see Schlenker & Pontari, 2000; Schlenker & Weigold, 1992). Second, self-serving biases are exacerbated under cognitive load (Paulhus, Graf, & Van Selst, 1989; Paulhus & Levitt, 1987) when people have difficulty engaging in strategic self-presentation (cf. Gilbert, Krull, & Pelham, 1988). Although both of these lines of research suggest that self-serving biases are indeed self-deceptive, they remain inconclusive as the results are sensitive to the same criticism: Both lines of research could be interpreted as evidence that self-serving biases are merely habitual or automatized (cf. Paulhus, 1993). The fact that self-serving biases actually increase under cognitive load may be evidence of an effortful correction process (Gilbert & Malone, 1995) that people learn to use to temper their self-serving statements as they mature, so that they are not seen as self-aggrandizing (an ineffective self-presentational strategy; Godfrey, Jones, & Lord, 1986; Paulhus, 1998).

The third class of evidence that SSBs are self-deceptive is that they have emerged even when people recognize that the accuracy of their statements can and will be checked. For example, college students inflated their SAT score even when they knew that their claims would be verified in the registrar database (Shepperd, 1993). This evidence suggests a self-deceptive component to self-serving biases, but it could be argued that people initially falsely report their SAT for self-presentational purposes, and over time this false report replaces their initial memory (e.g., Loftus, 1997). Thus, across the three types of evidence, it remains unclear whether self-serving biases are indeed self-deceptive.

Although support for the self-deceptive nature of self-serving biases may be inconclusive, the assumption that self-serving biases are self-deceptive is nevertheless widespread in social psychology. For example, self-serving biases are thought to be an important component of the positive illusions proposed to underlie mental and physical health (Taylor, 1989; Taylor & Brown, 1988), and are also thought to support the theories that link self-attributes to desirable future outcomes (Kunda, 1987).

The Nature of Self-Deception

Given the importance of the assumption that self-serving biases are self-deceptive, it may seem somewhat surprising that there has not been a more direct test of the hypothesis. In large part, however, the hypothesis remains only partially tested because of historical difficulties defining and operationalizing self-deception. Philosophers, clinicians, and scientists have all struggled with the paradox of self-deception - how is it that people could deceive themselves? Early philosophers, debating whether self-deception was theoretically possible, were more concerned about the characteristics of consciousness than a definition of self-deception (e.g., Fingarette, 1969). For example, in order for a person to be *able* to deceive himself (and not realize that he is doing so), consciousness must be nonunitary and nontransparent (Plato, 386 B.C./1953; Sartre, 1958/1943). This idea was later accepted by Freud (1923/1961), who expanded on it by arguing that awareness and nonawareness of particular bits of information can be motivated. The active repression of unwanted thoughts into the unconscious became the cornerstone of psychoanalysis, and set the stage for researchers to begin studying self-deception. Numerous researchers capitalized on Freud's ideas by arguing that this "motivated forgetting" is crucial to self-deception, and can explain why self-deception plays such an important role in behaviors such as completing self-report inventories (Anastasi, 1961; Meehl & Hathaway, 1946), defense mechanisms (Hilgard, 1949), problem-solving (Wason & Johnson-Laird, 1972), and perceptual defense (Murphy, 1970, 1975).

Despite all of the hypothesizing about the effects of self-deception, Sackeim and Gur (1978; Gur & Sackeim, 1979; Sackeim, 1988) made the first concerted effort to define *and* empirically demonstrate it. Building on the initial discussions of Plato, Freud, Fingarette, and others, Sackeim and Gur offered a definition of self-deception that incorporated many of the basic elements that had been discussed previously. They argued that four criteria were necessary and sufficient for ascribing self-deception to any given phenomenon – an individual must (1) simultaneously (2) hold two contradictory beliefs, (3) one of which the individual is not aware of (4) because of a motivated act. In this manner, self-deception was contrasted with other-deception, or intentionally deceiving someone other than the self, and was thought to occur when individuals are confronted with aspects of themselves that they find difficult to accept.

Gur and Sackeim (1979) used a voice recognition paradigm to demonstrate self-deception. Specifically, they had participants record voice samples that were then played back to them interspersed with voice samples from other people. Participants were then asked to determine whether each voice sample belonged to the self or not. Galvanic skin responses (GSR) to each voice were recorded. Because all participants showed a higher GSR to voice samples from the self, people's GSR could be compared to their verbal identification of the sample to see if any self-deception was occurring. If participants' GSR indicated that the voice had been identified as the self, yet they verbally reported that the sample did not come from the self, then the first two criteria for self-deception have been met (i.e., participants simultaneously said the voice sample was not theirs and their GSR indicated that, at some level, they knew it was).

To establish the third criterion for self-deception (i.e., the individual is not aware of the two contradictory beliefs), all participants went through an extensive post-experimental interview where they reported whether they were aware of having committed any errors. Very few participants were able to report whether or not they had made errors, and if so, what type they were (e.g., saying self to a voice sample when it was not self, saying not self to a voice sample that was self). The fourth criterion, the *motivated* unawareness of one of the contradictory beliefs, was the hardest to establish. Gur and Sackeim defined self-confrontation as correctly identifying a voice sample as the self. They hypothesized that people who are dispositionally or situationally averse to self-confrontation would be more likely to make errors in identifying a voice sample as the self. An individual difference measure of aversiveness to self-confrontation developed by Sackeim and Gur (called the Self-Deception Questionnaire; 1979) as well as a situational manipulation of failure both indicated this to be true; participant's threatened by self-confrontation were more likely to be self-deceptive.

It appears that Gur and Sackeim (1979) proposed a definition of self-deception that was testable, and their studies offered the first empirical evidence that self-deception does exist. However, their definition of self-deception was still paradoxical in the sense that a self-deceptive person had to know and not know something at the same time. Greenwald (1988) later offered a non-paradoxical definition of self-deception intended to make Sackeim and Gur's (1978) original conceptualization less restrictive. Rather than make the assumption that one must know unconsciously what is being avoided consciously, Greenwald defined self-deception as knowledge avoidance. He argued that cognitive analysis of incoming information happens in a series of stages. The first stage would be a crude analysis of the incoming information, whereas the second stage would be more elaborative processing. If information is discarded before it reaches the second stage, then no further processing would be able to occur. Thus, self-deception to Greenwald would be detecting any unwelcome content at the early stage of analysis and then failing to process that information further.

Paulhus (1991) also attempted to offer a nonparadoxical definition of self-deception that built on Sackeim and Gur's pioneering work. Specifically, he was interested in developing an individual difference measure to tap self-deceptive and other-deceptive tendencies. He reformulated the items from Sackeim and Gur's Self- and Other-Deception Questionnaires (1979) and called the resulting scale the Balanced Inventory of Desirable Responding (BIDR; 1991). The BIDR contains two subscales; one is thought to measure self-deception, whereas the other is thought to measure other-deception or impression management. To assess self-deceptive tendencies, respondents complete items that reflect the tendency to have an unrealistically (and presumably impossibly) positive view of the self. For example, a respondent who strongly endorses the item "I never regret my decisions" is presumably being self-deceptive because everyone regrets a decision at some point. Paulhus argues that people who score high on this subscale are self-deceptive because they honestly believe the unrealistically positive view of themselves that they report.

Reconceptualizing Self-Deception

Given the complexity of prior definitions of self-deception, it seems prudent to offer a simpler definition that is also easy to test. Sackeim and Gur (1978) offer a clinical definition where a person simultaneously holds two beliefs that are contradictory, one of which the person is unaware of because of a motivated act. While thorough, this definition is complex, restrictive, and difficult to test, as evidenced by Gur and Sackeim's (1979) empirical demonstration. Greenwald (1988) proposes a definition of self-deception in terms of knowledge avoidance, but his framework ignores large classes of behavior that would be regarded as self-deceptive by most people (such as occurs when people reinterpret or rationalize their negative behavior as positively motivated; see Steele, 1988). Likewise, Paulhus' (1991) definition of self-deception as an unrealistically positive view of the self is unnecessarily limited to the case of information about the self. Thus, in the current paper, we offer a more basic and broad definition of self-deception. At a fundamental level, self-deception can be defined as *misleading oneself so that the self benefits in some way*. Because the concept of “misleading” involves either deliberate deception, supplying incorrect information, or failing to supply correct information, this definition allows for both “hot” and “cold” forms of self-deception.

In the present research we propose a new measure of self-deception that reflects the broader definition we have proposed, and also follows the lead of Gur and Sackeim (1979) by relying on behavioral indicators rather than self-report. Specifically, we developed a technique that allows people to cheat in a manner that is ostensibly outside of the experimenter’s awareness, and under some circumstances, also ostensibly accidental. In order to accomplish this goal, we presented people with two series of math problems that were designed to be onerous but solvable (i.e., the problems would take some time to complete and would potentially be frustrating). In both series of problems people were told that there was a "bug" in the computer software, and that they would have to hit the spacebar as soon as each problem appeared on the screen in order to prevent the answer from being unintentionally displayed. It was stressed to participants the importance of hitting the spacebar in a timely fashion, as we would not know when they saw the answers and when they did not, and thus we would have no way of knowing if the experiment was conducted properly.

Participants were further told that in the first series of problems it would be easy to hit the space bar in time, and in fact, they had ten seconds to do so. For the second series of problems, however, participants were warned that it would be difficult to hit the space bar in time, and that they would have to react as fast as they could. In this series of trials they had only one second to hit the space bar. This is a sufficient amount of time for a task like hitting a computer key (e.g., see Meyer & Schvaneveldt, 1971), but the interval is short enough that it gives participants an excuse for not hitting the key in time. It is important to remember that participants were told that the experimenter would not know whether they had hit the key in time, and thus this excuse was intended primarily for purposes of self-justification rather than self-presentation. Thus, the self-deceptive aspect of this cheating was not whether they used the answer that was revealed to them, as all participants would presumably know they are relying on a provided answer, but rather whether they were able to mislead themselves about whether they intended the answer to be revealed in the first place.

In order to isolate self-deceptive cheating in this task, the goal of the analyses was to predict residual cheating on the second task after variance predicted by cheating on the first task (i.e., self-aware cheating) was removed. According to the definition of self-deception provided earlier, this measure of residual cheating is an index of self-deception because it allows us to isolate the tendency to cheat on problems that would be taxing and difficult for the self when one can make an excuse for one's behavior. Furthermore, this excuse is only directed toward the self, because the experimenter is ostensibly oblivious to any cheating that occurs in either the slow or fast conditions. Thus, residual cheating on the fast but not the slow math task can be taken as evidence of self-deception.

Relationship Between Self-Deception and the Self-Serving Bias

If the SSB is self-deceptive, then it should correlate with our new measure of self-deception. Of the many measures of self-serving bias available to test this possibility, we chose a measure that does not rely on a comparison of the individual to the population mean, as it is difficult to know whether any single person is really better than average in any particular domain. Such comparison measures reliably demonstrate self-serving biases in the aggregate, as not everyone can be better than average, but they cannot unambiguously show that a particular person is being self-serving. For this reason, in the current experiments we relied on a type of self-serving bias whereby people claim that tasks at which they have succeeded are inherently more important than tasks at which they have failed (Tesser & Paulhus, 1983). This measure does not involve claiming to be better than one is, but rather involves minimizing one's weaknesses and maximizing one's strengths. Although slightly different from the classic measure of SSB (attributing success to internal factors and failure to external factors), this type of measure is consistent with the logic of associating the self with desirable events or outcomes and distancing the self from undesirable events or outcomes. Such alternative measures of SSB have been used successfully in previous research (see Beauregard & Dunning, 2001; Dunning, Meyerowitz, & Holzberg, 1989; Shepperd, 1993).

In order to create an individual difference measure that relies on this type of self-serving bias, we presented people with two novel tasks and gave them success feedback on one and failure feedback on the other. The novelty of the tasks was important for two reasons: 1) it increased the likelihood that participants would believe the false feedback, and 2) it ensured that they would have no basis for knowing whether the tasks were important or trivial. In contrast to Tesser and Paulhus (1983), however, we did not give participants the impression that either task was particularly meaningful. Indeed, the tasks were downplayed somewhat by communicating that they were new and we were not exactly sure what they measured. Our goal in eliminating this aspect of Tesser and Paulhus' procedure was to decrease the overall magnitude of participants' self-serving bias, as participants should be less likely to be self-serving on a relatively meaningless task (see Tesser & Paulhus, 1983). With this omission it thereby seemed more likely that we would create a normally distributed measure of SSB that avoided ceiling effects. Self-serving bias was measured in this procedure as the degree to which people claimed (or did not claim) that the task at which they had succeeded tapped more important qualities than the task at which they had failed. The goal of the first experiment was to assess whether this measure of self-serving bias correlated with our measure of self-deception.

Experiment 1

Method

Participants. 108 students participated in partial fulfillment of course requirements for introductory psychology. Each participant completed the experiment in a private cubicle.

Procedure. Participants were told that we were in the process of developing several computerized tasks that were previously paper-and-pencil measures. The first task was a Mental Math Task, which consisted of 2 sets of 10 equations, with the two sets separated by a brief break. Each equation in each set contained 10 numbers (ranging from 1 to 20) that participants were to add or subtract. The problems were designed to be onerous and time-consuming (i.e., completing them would be aversive), but due to the simple mathematics that were involved, they were not inherently difficult to solve. Participants were told that if they answered a problem incorrectly, it would re-appear on the screen until they provided the correct answer. Thus, all participants had to provide a correct answer to all 20 questions in order to complete the Mental Math Task.

Participants were directed to hit the spacebar immediately upon presentation of each equation so that a response box would appear on the screen, and were told that as soon as the response box appeared they would be given as much time as they needed to solve the equation. They were specifically told that they should not try to solve the problem prior to hitting the space bar, but rather they were to hit the space bar first, thereby causing the response box to appear, and then they were to solve the problem. They were also informed that in this early stage of program development there was still a "bug" in the software, whereby if they took too long to hit the spacebar, the correct answer would unintentionally appear on the screen. Participants were told that for the first set of problems (which we refer to as the "slow math set"), this software bug was not a major problem as they had plenty of time to hit the spacebar before the correct answer would appear (the actual delay was 10 seconds). For the second set of problems (which we refer to as the "fast math set"), however, participants were warned that the software bug was more difficult to avoid, as the correct answer would appear if the spacebar was not hit within one second after the equation appeared. The experimenter reiterated that such a small amount of time may make it difficult for some people to hit the spacebar quickly enough. The experimenter also noted that it was important to try to hit the spacebar as rapidly as possible, as we would have no way of knowing whether the software bug disrupted the experimental procedure, thereby invalidating their data.

After completing the Mental Math Task, participants were told that the next tasks were being conducted in conjunction with a local mental hospital trying to understand how accurate diagnoses are made. In a counterbalanced order, participants were presented with a Suicide Note Task and a Schizophrenic Definitions Task. Each task was described in very abstract terms – the qualities and traits measured by each task were not delineated, nor was the importance of doing well on either task emphasized. For the Suicide Note Task participants read a series of 26 suicide notes, some of which were real and some of which were fake (Ross, Lepper, & Hubbard, 1975). After reading each note participants were asked to indicate whether the note was real or fake. For the Schizophrenic Definitions Task participants read a series of 26 word definitions that displayed moderately disordered thinking (Manis & Paskewitz, 1984). After reading each definition, participants were asked to indicate whether it was provided by a patient who was schizophrenic or not. Upon completion of these tasks participants received false feedback about their performance such that, in a counterbalanced fashion, they were told they did well on one task (top 17%) and poorly on the other (bottom 28%). After participants received this feedback they were asked to rate how important the skills are that are necessary to complete each task successfully. Responses to these two questions were made on a seven-point scale, ranging from "Very Unimportant" (1) to "Very Important" (7).

Participants were then thoroughly probed for suspicion. First, they were asked whether anything seemed strange or unusual in the experiment. Next they were asked what they thought the experiment was attempting to demonstrate. Last they were told that the experiment wasn't really concerned with validating computerized tasks and were asked to guess what the true goal might have been. Although some participants in all three experiments voiced suspicions at this third stage, none raised the issue of the math tasks and the unintentional display of answers. The only systematic concerns that were raised focused on the accuracy of the feedback on the two novel tasks. These participants wondered if perhaps we were trying to make them feel bad or good via the feedback. Including or eliminating these participants from the analyses had no impact on the results in any of the three experiments, and thus all participants were included.

Results

The first question was whether participants cheated on the slow and fast math tasks. Consistent with the onerous nature of the tasks, and with the verisimilitude of our claim that we would not be able to ascertain whether participants had seen the correct answers, substantial cheating occurred on both math tasks. Of the 10 problems in each series, participants waited for the answer to appear prior to hitting the space bar on an average of 3.97 problems in the slow math set and 4.93 problems in the fast math set, indicating significant cheating in both conditions, p's $< .001$, and greater cheating in the fast math set than the slow math set, $F(1,107) = 14.78$, $p < .001$.

The next question was whether participants, on average, displayed a self-serving bias when rating the importance of the tasks at which they had succeeded and failed. Because we did not emphasize the importance of the two tasks, it was unclear whether most participants would be self-serving. Ratings of the importance of the two tasks were subjected to a two (task) by two (feedback) within subjects ANOVA. This analysis did not reveal an interaction between task and feedback, $F < 1$, ns, suggesting that no overall self-serving bias emerged on this measure. In order to create an idiographic measure of self-serving bias, a difference score was then computed whereby the rated importance of the task at which participants had ostensibly failed was subtracted from the rated importance of the task at which they had ostensibly succeeded. Consistent with the absence of a significant interaction, this difference score had a mean of -.03 (suicide notes task, M = 4.51, SD = 1.68; word definitions task, M = 4.61, SD = 1.39). Thus, it appears that although people rated both tasks of moderate importance, there was certainly no ceiling effect on the self-serving bias measure, nor did most of the participants show a self-serving pattern of importance ratings.

The final, and central, question in the analysis was whether the extent to which participants showed a self-serving bias predicted the degree to which they cheated on the fast but not the slow math task. In order to address this question, a regression analysis was conducted in which cheating on the fast math task was the dependent variable and cheating on the slow math task was entered as the first predictor, followed by the differential importance score (i.e., the self-serving bias measure) as the second predictor. The goal of this analysis was to partial out the degree to which participants were knowingly cheating (which they could obviously do on both the slow and fast math tasks, but presumably accounted for nearly 100% of the cheating on the slow math task), and assess whether differential importance predicted residual variance in cheating on the fast math task beyond that predicted by the slow math task. Thus, this analysis was intended to remove the intentional cheating component from the fast math task, leaving only error variance and self-deception as the basis for the residual variance in cheating. Consistent with predictions, differential importance added significant variance to the prediction of cheating on the fast math task ($\beta = .14$, $R^2\ \Delta = .02$, $F(1,106) = 4.71$, $p < .04$) beyond that predicted by cheating on the slow math task ($\beta = .71$, $R^2\ \Delta = .49$, $F(1,107) = 116.77$, $p < .001$). [1] Furthermore, and importantly, this regression equation was not reversible. That is, when cheating on the fast math task and differential importance were used to predict cheating on the slow math task, differential importance did not predict residual variance in cheating on the slow math task (β = -.03, *p* > .65) beyond that predicted by cheating on the fast math task (β = .74, *p* < .001).

Although this analytic strategy of focusing on residual cheating demonstrates most clearly the relationship between self-serving biases and self-deception, it has the disadvantage of not indicating whether differences in cheating emerge between those who show a self-serving bias and those who do not. Thus, as a secondary analytic strategy, we also compared levels of cheating in both the slow and fast math tasks between those who were self-serving (difference score > 0; N = 32) and those who were not (difference score ≤ 0; N = 76). Despite the rather imprecise nature of this analysis, it revealed that self-serving participants showed greater cheating on the fast math task (M = 6.31, SD = 3.32) than non-self-serving participants (M = 4.34, SD = 3.46), $t(106) = 2.74$, $p < .01$. No reliable differences emerged in cheating on the slow math task between self-serving (M = 4.62, SD = 3.43) and non-self-serving participants (M = 3.70, SD = 3.58), $t(106) = 1.24$, $p > .20$. This analysis is rather imprecise, as not all cheating on the fast math task is self-deceptive, but it nonetheless provides additional evidence that people who show a SSB are also likely to be self-deceptive.

Discussion

The results of Experiment 1 suggest that self-serving biases are self-deceptive. In particular, the tendency to rate a task at which one succeeded as more important than a task at which one failed predicted the tendency to cheat on a math task when that cheating could be justified as accidental, but not when that cheating was obviously purposeful. Although the size of the relationship between SSB and self-deception was quite small, given the large discrepancy between the nature of the two tasks (cheating on a math task vs. aggrandizing success/

minimizing failure on the suicide notes/word definitions tasks), the emergence of such a relationship is noteworthy nonetheless.

It is worth pointing out that the current experiment did not reveal any evidence for an overall self-serving bias across participants. As mentioned earlier, it seems likely that a robust self-serving bias did not emerge on the differential importance measure because the meaningfulness of the two novel tasks was never emphasized to participants, and thus many of them probably felt little need to stress the importance of their success over their failure. For the purposes of the current research, however, the absence of a strong self-serving bias may even be beneficial, as it increases the variance and thus the utility of the idiographic measure of self-serving bias. Even more importantly, because the self-deception measure involved a task that was not terribly important or ego-involving (i.e., the self-deception simply enabled participants to save a little time and effort), it may be the case that the relatively equivalent (in this case low) importance of each task is necessary in order for a relationship between them to emerge. If the self-serving bias task were designed to be of such high importance that most participants were self-serving, and yet the self-deception task was of such low importance that most participants were not self-deceptive, then their incommensurate importance levels might reduce the likelihood that a relationship would emerge between them. For this reason, no effort was made in subsequent experiments to increase the perceived meaningfulness of the two novel tasks and the consequent likelihood that most participants would show a self-serving bias.

Experiment 2

The findings from Experiment 1 were as predicted, but one issue that was not addressed was whether self-serving bias predicts variance in self-deception beyond that predicted by a self-report measure. Despite the fact that the self-report measures of self-deception were designed to fit a different definition of self-deception (one focused on self-enhancement), it was deemed worthwhile to demonstrate that the measure of SSB would predict variance in self-deceptive behavior beyond that predicted by a self-report measure of self-deception. Thus, the goals of Experiment 2 were to assess whether a self-report measure of self-deception also predicts unique variance in our measure of self-deceptive cheating, and to assess whether self-serving bias predicts unique variance in self-deception beyond that predicted by a self-report scale.

Method

Participants. Fifty-two students participated in partial fulfillment of course requirements for introductory psychology.

Procedure. The materials and procedure of Experiment 2 were identical to that of Experiment 1, with the exception that after completing the self-deception and self-serving bias tasks, participants then completed the 20 item self-deception subscale of the Balanced Inventory of Desirable Responding (BIDR-SD; Paulhus, 1991).

Results

Analyses again revealed significant cheating on both the slow and fast math tasks (p's < .001), and again greater cheating on the fast (4.14) than the slow task (3.10), $F(1,51) = 12.33$, $p < .001$.

Consistent with Experiment 1, participants in this experiment did not show evidence of a self-serving bias (interaction between task and feedback, $F < 1$, ns). A differential importance score was computed as in Experiment 1, with a mean of .06 (suicide notes task, M = 4.27, SD = 1.69; word definitions task, M = 4.56, SD = 1.09).

The central question of this experiment was whether self-serving bias predicted independent variance in self-deceptive cheating, beyond that predicted by the BIDR-SD. In order to address this question, a regression analysis was conducted in which cheating on the fast math task was the dependent variable and cheating on the slow math task was entered as the first predictor, followed by the self-deception subscale, followed by the differential importance measure. Consistent with predictions, differential importance added significant variance to the prediction of cheating on the fast math task ($\beta = .19$, $R^2 \Delta = .03$, $F(1,48) = 5.21$, $p < .03$) beyond that predicted by cheating on the slow math task ($\beta = .76$, $R^2 = .65$, $F(1,50) = 91.54$, $p < .001$) and the BIDR-SD ($\beta = .11$, $R^2 \Delta = .02$, $F(1,49) = 2.71$, $p > .10$).[2] As in Experiment 1, this regression equation was not reversible. When cheating on the fast math task and differential importance were used to predict cheating on the slow math task, differential importance did not predict residual variance in cheating on the slow math task ($\beta = -.09$, $p > .30$) beyond that predicted by cheating on the fast math task ($\beta = .85$, $p < .001$) and the BIDR-SD ($\beta = -.07$, $p > .40$). Correlational analyses further revealed that scores on the BIDR-SD were not correlated with cheating on either task or with the differential importance score (r's $\leq .21$, p's $> .10$).

Finally, as in Experiment 1, we also compared levels of cheating in both the slow and fast math tasks between those who were self-serving (N = 18) and those who were not (N = 34). Consistent with Experiment 1, this analysis revealed that self-serving participants showed greater cheating on the fast math task (M = 5.61, SD = 3.28) than non-self-serving participants (M = 3.35, SD = 3.58), $t(50) = 2.27$, $p < .03$. A marginally significant difference also emerged in cheating on the slow math task between self-serving (M = 4.11, SD = 3.61) and non-self-serving participants (M = 2.56, SD = 2.79), $t(50) = 1.72$, $p < .10$.

Discussion

The results of Experiment 2 replicate those of Experiment 1, in that self-serving bias was reliably related to self-deception. Furthermore, the results of Experiment 2 extend those of

Experiment 1 by demonstrating that self-serving bias predicts self-deception beyond the degree to which it is predicted by a widely used self-report measure, the BIDR-SD. Indeed, the BIDR-SD failed to predict self-deceptive cheating in this experiment. The fact that the BIDR-SD has been shown to predict other behaviors that appear self-deceptive (e.g., positive reappraisal, illusion of control, overconfidence; Paulhus, 1991) suggests that such self-enhancing self-deceptive behavior may be more readily predicted by this self-report inventory than the type of self-deception involved in the cheating task.

Although we have argued that the lack of a relationship between the BIDR-SD and the residual cheating measure suggests that the BIDR-SD is not well suited for measuring the type of self-deception tapped by this measure, it could also be argued that the lack of a relationship suggests that the residual cheating measure itself is a poor measure of self-deception. Thus, it seems worthwhile to address concerns about the validity of the cheating measure as a measure of self-deception. According to the logic presented earlier, we are suggesting that in addition to self-aware cheating, some self-deceptive cheating was also occurring on the fast mental math task, which explains why cheating rates were higher in the fast set than in the slow set. Recall that participants completed the experiment alone in a cubicle and thought that the experimenter would have no idea how many of the answers were actually seen. Thus, it makes sense that some of the increase in cheating from the slow task to the fast task represents self-deception: people cheat more because they mislead themselves into believing that they were unable to hit the spacebar within the one-second time frame that would have prevented the answer from appearing. The untested crux of this argument, however, is that people actually have the ability to hit the space bar within the one-second time frame in the context of this experiment. Although ample previous research has demonstrated that people are capable of making complex cognitive responses within one second, a demonstration of this ability within the present paradigm is also important. In addition, demonstrating that people are able to hit the spacebar within one second when it is advantageous to them to do so would strengthen the claim that the cheating measure really does measure self-deception.

In order to accomplish this goal, Experiment 3 contained an additional series of math tasks in which participants were required to hit the key quickly in order to prevent the math problems from becoming more difficult. According to our predictions, the measure of self-serving bias should predict inability to respond within one second when the answer is thereby caused to appear, but not inability to respond within one second when the problem is thereby made more difficult.

Inclusion of this measure also addresses an alternative explanation for the findings of Experiments 1 and 2. It could be argued that there is something unique about the people who showed a self-serving bias that prevents them from hitting the spacebar within one second. That is, maybe the reason that the self-serving bias was correlated with cheating on the second task had nothing to do with self-deception, but was just driven by the inability of self-serving participants to respond quickly enough. The inclusion of this new measure can rule out this possibility.

Another alternative explanation for the results of Experiments 1 and 2 is that the relationship between self-serving bias and self-deception is actually a spurious correlation driven by some participants thinking too much. That is, it is possible that participants who appear to be self-serving are simply high in Need for Cognition, and thus they thoughtfully compared their current success and failure to their past performance and determined that the tasks at which they succeed tend to be more important than the tasks at which they fail. These same thoughtful individuals might then forget to hit the spacebar when confronted with the math problems in the fast set because they get absorbed in solving the problem, and forget to hit the space bar to prevent the answer from appearing. In order to rule out this possibility, the Need for Cognition scale (Cacioppo & Petty, 1982) was included in Experiment 3. According to our predictions, Need for Cognition should be unrelated to the idiographic measure of self-serving bias, and if any relationship emerges between Need for Cognition and cheating it should be a negative one (i.e., people high in Need for Cognition should avoid cheating as they enjoy solving problems).

Experiment 3

Method

Participants. One hundred forty three students participated in partial fulfillment of course requirements for introductory psychology.

Procedure. The materials and procedure of Experiment 3 were identical to that of Experiment 2 up through the administration of the BIDR-SD. Following this scale, participants completed the Need for Cognition scale and then a second series of math equations. For this second series of questions participants were told that the task would combine mental and physical prowess by requiring them to hit a key very quickly and then solve an equation. Participants were told that they would encounter a series of equations that consist of four numbers each, but if they took too long to hit the spacebar to initiate the response box, this equation would be replaced by an equation that contained ten numbers. As with the first math tasks, these equations were split into two sets of 10 equations. In Set 1 participants were allowed 10 seconds to hit the spacebar without consequence, whereas in Set 2 they were allowed only 1 second to hit the spacebar. Because it might have seemed odd to give participants such a long time to respond in an ostensible measure of physical prowess, participants were told that Set 1 of this task would be easy and was intended as a warm-up for Set 2, where they would only have 1 second to hit the spacebar before their 4-number equation was replaced by a 10-number equation. They were also told that after hitting the key they would have as much time as they needed to calculate the solution.

Results and Discussion

Consistent with Experiments 1 and 2, no evidence emerged for a self-serving bias (interaction between task and feedback, $F < 1$, ns). Differential importance was computed as in Experiments 1 and 2, with a mean of .06 (suicide notes task, M = 4.63, SD = 1.54; word definitions task, M = 4.64, SD = 1.29).

Analyses again revealed significant cheating on both the slow and fast math tests (p's < .001), and again greater cheating on the fast (3.59) than the slow test (2.39), $F(1,142) = 40.29$, $p < .001$. Despite this relative inability to hit the spacebar in a timely fashion during the first two math series, participants nevertheless responded more quickly during the second two math series when failure to hit the space bar caused the problems to become more difficult. In the second series participants only failed to hit the space bar an average of .1 times in the slow test and .8 times in the fast test. Despite their small size, these means were significantly different from zero (p's < .01), and also different from each other, $F(1,132) = 26.51$, $p < .001$.[3] This finding suggests that although it was at least somewhat difficult for participants to always hit the space bar quickly enough, one second is adequate time to make this response for most people most of the time.

Because failure to hit the space bar in a timely fashion in the second series of math equations made the problems more difficult, this failure rate can be interpreted as each participant's baseline error rate, or inability to hit the space bar within the necessary time parameters. By subtracting each participant's error rates from their cheating rates, this measure allows a test of the hypothesis that the increase in cheating from the slow to the fast math task represents the addition of self-deceptive cheating and not just the addition of error. In order to test this possibility, failures to hit the space bar in the second slow and fast math tasks (which caused the problems to become more difficult) were subtracted from failures to hit the space bar in the first slow and fast math tasks (which caused the answers to appear), respectively. These difference scores were then compared, and they were found to differ marginally from one another, $t(132) = 1.88$, $p = .062$.

As in Experiment 2, the BIDR-SD was not correlated with differential importance nor with cheating on either the slow or fast math task, nor did it predict allowing the problems to become more difficult in either the slow or fast math task (r's $\leq .14$, p's $> .10$). Need for Cognition was not associated with the differential importance score ($r = -.01$, ns), and it was correlated positively with the self-deception subscale ($r = .20$, $p < .03$). As predicted, Need for Cognition was correlated negatively with cheating on both the slow ($r = -.22$, $p < .01$) and fast ($r = -.23$, $p < .01$) math tasks, and was unrelated to allowing the problems to become more difficult in both the slow ($r = -.13$, $p > .10$) and fast math tasks ($r = -.01$, ns).

Most importantly, regression analyses again revealed that differential importance predicted significant variance in cheating on the fast math task ($\beta = .12$, $R^2\ \Delta = .02$, $F(1,123) = 4.12$, $p < .05$) beyond that predicted by cheating on the slow math task ($\beta = .75$, $R^2 = .55$, $F(1,126) = 153.20$, $p < .001$), the self-deception subscale ($\beta = .00$, $R^2\ \Delta = .00$, $F(1,125) = 0.00$, ns), and Need for Cognition ($\beta = -.06$, $R^2\ \Delta = .00$, $F(1,124) = 1.07$, $p > .30$).[4] Additionally, when the variables of allowing the problems to become more difficult in the slow and fast tasks were added to this equation as predictor variables, neither accounted for significant variance in cheating on the fast math task, and differential importance remained a marginally reliable predictor, (β = .12, p = .052). These results emerged despite reliable bivariate correlations between failure to hit the space bar in time on the fast math task and failure to hit it in time on the slow (r = .22, $p < .05$) and fast (r = .26, $p < .01$) tasks in which the problems became more difficult. This finding suggests that the fast math task is indeed tapping differential cheating rates, and not simply an inability to hit the space bar within the one-second time constraints.

As in Experiments 1 and 2, these regression equations were not reversible. When cheating on the fast math task and differential importance were used to predict cheating on the slow math task, differential importance actually predicted a significant *decrease* in residual cheating on the slow math task (β = -.17, $p < .01$) beyond that predicted by cheating on the fast math task (β = .73, $p < .001$), the self-deception subscale (β = -.01, $p > .85$), and Need for Cognition (β = -.06, $p > .30$). When this regression analysis was conducted with the dependent variable of allowing the problem to become more difficult in the fast math task, the only significant predictor was allowing the problem to become more difficult in the slow math task, (β = .19, R^2 = .04, $F(1,126) = 4.55$, $p < .04$).

Finally, as in Experiments 1 and 2, we also compared levels of cheating in both the slow and fast math tasks between those who were self-serving (N = 48) and those who were not (N = 95). Inconsistent with Experiments 1 and 2, this analysis failed to reveal differences in cheating on the fast math task between self-serving participants (M = 3.44, SD = 3.23) and non-self-serving participants (M = 3.67, SD = 3.32), $t(141) = .40$, $p > .65$. Similarly, no differences emerged in cheating on the slow math task between self-serving participants (M = 2.04, SD = 2.65) and non-self-serving participants (M = 2.56, SD = 2.85), $t(141) = 1.05$, $p > .25$. This failure to replicate the prior experiments despite the consistency in the prediction of residual cheating may reflect the insensitive nature of this test. Alternatively, this failure to replicate may implicate a suppression effect, which is suggested by the presence of a partial correlation between differential importance and cheating on the fast task in the absence of a bivariate relationship between these variables.

In sum, the results of Experiment 3 are consistent with the findings from Experiments 1 and 2, in that self-serving bias again predicted residual cheating on the fast math task. Because Need for Cognition was not correlated with our self-serving bias measure, and was negatively correlated with cheating, these findings suggest that the "thoughtful" reinterpretation of people's responses on the self-serving bias measure is unlikely. This study also provides an important addition to the earlier findings by demonstrating that when failure to respond quickly hurt rather than helped participants, the relationship between self-serving bias and task performance melted away. In so doing, these results suggest that the earlier findings were not an artifact of an inability on the part of some participants to respond quickly enough. Indeed, the results of Experiment 3 indicate that when baseline error rates were removed from responses, self-serving participants still showed an increase in cheating in the task that allowed self-deception. It is important to note that this increase in cheating was only marginally significant, but it is also the case that this procedure provided a particularly conservative test of the hypothesis. That is, in this procedure participants who were able to deceive themselves into failing to respond in a timely fashion in the first math task were then confronted with an almost identical math task in which self-interest now pushed them into quick responding. It is possible that the similarity in circumstances may have forced at least some participants to maintain their self-deception, requiring them to show internal consistency in their response patterns and thereby necessitating an occasional error in order to sustain the fiction that they were unable to react quickly enough in the earlier task. Of course, this possibility is entirely conjectural, and future research will be necessary to sort it out.

General Discussion

The results of three experiments are consistent with the prediction that self-serving biases have a self-deceptive component. Across three studies, people who rated a verbal task at which they succeeded as more important than a verbal task at which they failed tended to cheat on a series of math problems. Because this relationship between self-serving biases and cheating only emerged when people could justify the cheating to themselves, these findings highlight the importance of self-deception in this process. Furthermore, these effects emerged despite the fact that a self-report measure of self-deception (i.e., the self-deception subscale of the BIDR) did not predict the tendency to cheat on either the self-aware or the self-deceptive cheating tasks. Thus, these results provide evidence from procedures highly discrepant from those used in previous research (e.g., Paulhus et al., 1989; Schlenker & Pontari, 2000; Shepperd, 1993) that there is a self-deceptive component to self-serving biases.

Perhaps the most important aspect of these experiments is not what they tell us about self-serving biases, as to a great degree the results are consistent with the zeitgeist in the field. Rather, what is more novel about the current results is that they provide evidence concerning the empirical tractability of a more inclusive definition of self-deception, and of a procedure designed to capture it. Given the private nature of the Mental Math Task in the current studies, there was no reason for participants not to cheat to a similar degree on the slow math task as they did on the fast math task. After all, participants worked independently in private rooms, and they were told that the experimenter would not know whether they had seen any of the answers. Testifying to their belief in this claim is the substantial cheating that emerged in both the slow and fast math tasks. Thus, the most plausible explanation for the differential cheating is that participants justified waiting for the answers in the fast condition by telling themselves that they simply could not hit the spacebar in time to avoid seeing the answer. The results from the task in which the problems become more difficult testify that this is indeed a deception, as these same participants were quite capable of hitting the space bar when it facilitated rather than interfered with their goals. Furthermore, this presumed inability to respond appears to have been a deception provided for the benefit of the self, as the experimenter was ostensibly unaware of what was happening in the participant's cubicle and participants never offered this justification to anyone else. In combination, these results suggest that participants misled themselves because it was beneficial for them to do so. They did not want to complete the onerous math problems, so their "inability" to hit the spacebar in a timely fashion allowed them to avoid the difficult task without feeling negative affect such as guilt.

It is also worth addressing why the self-deception subscale of the BIDR did not correlate with our new measure of self-deception. As noted earlier, this finding is not too surprising, given that the BIDR-SD and the cheating measure of self-deception are derived from fundamentally different definitions of self-deception. Paulhus (1991) defines self-deception as an honestly held, unrealistically positive view of the self. There is no reason to expect a scale derived from such a definition of self-deception to correlate with the cheating measure we developed here, which is based on a definition of self-deception as misleading the self. Nevertheless, one might argue that even constructs that are operationalized differently but are purportedly measures of the same latent variable should be correlated, even if that correlation is small. Given that this was not the case in Experiments 2 or 3, it is possible that the BIDR-SD is well suited to measure tendencies to hold an inflated view of self (e.g., overconfidence, illusions of control, positive reappraisal; Paulhus, 1991), but not self-deceptive tendencies that manifest themselves in other ways.

Finally, it should be noted that although these results provide an additional piece of the puzzle regarding self-deception and self-serving biases, there are obvious limitations to these data. First and foremost, it is difficult to establish with certainty that participants were indeed deceiving themselves in the Mental Math Task. The logic of the task suggests that they were, but there is no definitive marker of self-deceptive behavior, and thus the measure developed here awaits further validation. Second, although the SSB and self-deception were correlated in these experiments, the presence of such a relationship does not necessarily mean that the SSB has a self-deceptive component. It may simply be the case that people who are self-serving are also self-deceptive, as both of these tendencies may be driven by a third variable. For example, it could be the case that the joint presence of high explicit and low implicit self-esteem leads people both to knowingly self-present as confident and competent (SSB) and also to engage in a variety of self-deceptive behaviors that minimize self-doubt or maximize utility (cf. Jordan, Spencer, Zanna, Hoshino-Browne, & Correll, in press). Our goal in presenting this research has not been to close the door on all possible alternatives, but rather to open up the interesting, albeit thorny, problem of self-deception to an additional avenue of inquiry.

References

Anastasi, A. (1961). *Psychological testing*. New York: Macmillan.

Arkin, R. M., & Lakin, J. L. (2001). The Taj Mahal of selves. *Psychological Inquiry, 12*, 203-205.

Baumeister, R. (1998). The self. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (4th ed., pp. 680-740). New York: Oxford University Press.

Beauregard, K. S. & Dunning, D. (2001). Defining self-worth: Trait self-esteem moderates the use of self-serving trait definitions in social judgment. *Motivation and Emotion, 25*, 135-161.

Cacioppo, J. T., & Petty, R. E. (1982). The need for cognition. *Journal of Personality and Social Psychology, 42*, 116-131.

Dunning, D., Meyerowitz, J. A., & Holzberg, A. D. (1989). Ambiguity and self-evaluation: The role of idiosyncratic trait definitions in self-serving assessments of ability. *Journal of Personality & Social Psychology, 57*, 1082-1090.

Fazio, R. H., & Olson, M. A. (in press). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology*.

Fingarette, H. (1969). *Self-deception*. London: Routledge & Kegan Paul.

Freud, S. (1961). The ego and the id. In J. Strachey (Ed. & Trans.), *The standard edition of the complete psychological works of Sigmund Freud* (Vol. 19, pp. 3-66). London: Hogarth Press. (Original work published 1923).

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

Gilbert, D. T., Krull, D. S., & Pelham, B. W. (1988). Of thoughts unspoken: Social inference and the self-regulation of behavior. *Journal of Personality & Social Psychology, 55*, 685-694.

Godfrey, D. K., Jones, E. E., & Lord, C. G. (1986). Self-promotion is not ingratiating. *Journal of Personality and Social Psychology, 50*, 106-115.

Greenwald, A. G. (1988). Self-knowledge and self-deception. In J. S. Lockard & D. L. Paulhus (Eds.), *Self-deception: An adaptive mechanism?* (pp. 113-131). Englewood Cliffs, NJ: Prentice Hall.

Gur, R. C. & Sackeim, H. A. (1979). Self-deception: A concept in search of a phenomenon. *Journal of Personality and Social Psychology, 37*, 147-169.

Hilgard, E. R. (1949). Human motives and the concept of the self. *American Psychologist, 4*, 374-382.

Jordan, C. H., Spencer, S. J., Zanna, M. P., Hoshino-Browne, E., & Correll, J. (in press). Secure and defensive high self-esteem. Journal of Personality and Social Psychology.

Kunda, Z. (1987). Motivated inference: Self-serving generation and evaluation of causal theories. *Journal of Personality and Social Psychology, 53*, 636-647.

Loftus, E. F. (1997). Memories for a past that never was. *Current Directions in Psychological Science, 6*, 60-65.

Manis, M., & Paskewitz, J. (1984). Judging psychopathology: Expectation and contrast. *Journal of Experimental Social Psychology, 20*, 363-381.

Meehl, P. E., & Hathaway, S. R. (1946). The K factor as a suppressor variable in the Minnesota Multiphasic Personality Inventory. *Journal of Applied Psychology, 30*, 525-564.

Meyer, E. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology, 90*, 227-234.

Miller, D. T., & Ross, M. (1975). Self-serving biases in the attribution of causality: Fact or Fiction? *Psychological Bulletin, 82*, 231-225.

Murphy, G. (1970). Experiments in overcoming self-deception. *Psychophysiology, 6*, 790-799.

Murphy, G. (1975). *Outgrowing self-deception*. New York: Basic Books.

Paulhus, D. L. (1991). Measurement and control of response bias. In J. P. Robinson, P. R. Shaver, & L. S. Wrightsman (Eds.), *Measures of personality and social psychological attitudes* (pp. 17-59). San Diego: Academic Press.

Paulhus, D. L. (1993). Bypassing the will: The automatization of affirmations. In D. Wegner & J. W. Pennebaker (Eds.), *Handbook of mental control* (pp. 573-587). Englewood Cliffs, NJ: Prentice Hall.

Paulhus, D. L. (1998). Interpersonal and intrapsychic adaptiveness of trait self-enhancement: A mixed blessing? *Journal of Personality and Social Psychology, 74*, 1197-1208.

Paulhus, D. L., Graf, P., & Van Selst, M. (1989). Attentional load increases the positivity of self-presentation. *Social Cognition, 7*, 389-400.

Paulhus, D. L., & Levitt, K. (1987). Desirable responding triggered by affect: Automatic egotism? *Journal of Personality and Social Psychology, 52*, 245-259.

Plato. (1953). In R. Jowett (Ed. and Trans.), *The dialogues of Plato* (Vol. 3). Oxford: Clarendon Press. (Original work published circa 386 B.C.).

Ross, L., Lepper, M. R., & Hubbard, M. (1975). Perseverance in self-perception and social perception: Biased attributional processes in the debriefing paradigm. *Journal of Personality and Social Psychology, 32*, 880-892.

Sackeim, H. A. (1988). Self-deception: A synthesis. In J. S. Lockard & D. L. Paulhus (Eds.), *Self-deception: An adaptive mechanism?* (pp. 146-165). Englewood Cliffs, NJ: Prentice Hall.

Sackeim, H. A., & Gur, R. C. (1978). Self-deception, self-confrontation, and consciousness. In G. Schwartz & D. Shapiro (Eds.), *Consciousness and self-regulation: Advances in research and theory* (Vol. 2, pp. 139-197). New York: Plenum Press.

Sartre, J. P. (1958). *Being and nothingness: An essay on phenomenological ontology* (H. Barnes, Trans.). London: Methuen. (Original work published 1943).

Schlenker, B. R. (1980). *Impression management: The self-concept, social identity, and interpersonal relations*. Monterey, CA: Brooks/Cole.

Schlenker, B. R., & Pontari, B. A. (2000). The strategic control of information: Impression management and self-presentation in daily life. In A. Tesser & R. B. Felson (Eds.), *Psychological perspectives on self and identity* (pp. 199-232). Washington, DC: American Psychological Association.

Schlenker, B. R., & Weigold, M. F. (1992). Interpersonal processes involving impression regulation and management. *Annual Review of Psychology, 43*, 133-168.

Sedikides, C. (1993). Assessment, enhancement, and verification determinants of the self-evaluation process. *Journal of Personality and Social Psychology, 65*, 317-338.

Shepperd, J. (1993). Student derogation of the Scholastic Aptitude Test: Biases in perceptions and presentations of college board scores. *Basic and Applied Social Psychology, 14*, 455-473.

Steele, C.M. (1988). The psychology of self-affirmation: Sustaining the integrity of the self. *Advances in Experimental Social Psychology*, *21*, 261-302.

Taylor, S. E. (1989). *Positive illusions: Creative self-deceptions and the healthy mind*. New York: Basic Books.

Taylor, S. E., & Brown, J. D. (1988). Illusion and well-being: A social psychological perspective on mental health. *Psychological Bulletin, 103*, 193-210.

Tesser, A., & Paulhus, D. L. (1983). The definition of self: Private and public self-evaluation management strategies. *Journal of Personality and Social Psychology, 44*, 672-682.

Wason, P. C., & Johnson-Laird, P. N. (1972). *Psychology of reasoning: Structure and content*. Cambridge, MA: Harvard University Press.

Weary-Bradley, G. (1978). Self-serving biases in the attribution process: A reexamination of the fact or fiction question. *Journal of Personality and Social Psychology, 36*, 56-71.

Zuckerman, M. (1979). Attribution of success and failure revisited, or: The motivational bias is alive and well in attribution theory. *Journal of Personality, 47*, 245-287.

Endnotes

[1] Differential importance showed a marginally significant correlation with cheating on the slow task, $r = .16$, $p < .10$ and a significant correlation with cheating on the fast task, $r = .26$, $p < .01$. Cheating rates on the slow and fast tasks were correlated, $r = .73$, $p < .001$.

[2] Differential importance was not significantly correlated with cheating on the slow task, $r = .19$, $p > .15$, but was correlated with cheating on the fast task, $r = .35$, $p < .02$. Cheating rates on the slow and fast tasks were correlated, $r = .80$, $p < .001$.

[3] 10 participants failed to complete this final task for assorted reasons (e.g., they arrived late and the session was over by this point, they reported being late for class, or they took more than an hour to complete the experiment).

[4] Surprisingly, differential importance was negatively correlated with cheating on the slow task, $r = -.17$, $p < .05$ and uncorrelated with cheating on the fast task, $r = .00$, $p > .95$. Cheating rates on the slow and fast tasks were correlated, $r = .73$, $p < .001$.