CONCEPT ASSOCIATIONS AS THE BASIS OF SOCIAL PRIMING

BY

ERIKA SALOMON

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Psychology
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2016

Urbana, Illinois

Doctoral Committee:

Associate Professor Andrei Cimpian, Chair
Assistant Professor Michael W. Kraus, Yale University, Director of Research
Doctor Jesse L. Preston, University of California, Berkeley
Professor R. Chris Fraley
Professor Brent Roberts

ABSTRACT

Social priming has been a fundamental methodological practice in experimental social psychology for decades; however recent failures to replicate priming studies have cast some doubt on the robustness of subtle primes' influence on social behavior. The present research attempts to resolve this uncertainty by advancing the perspective that the robustness of social priming effects can be predicted a priori from the strength of the concept association under study. Across eight studies, reported without publication bias, I found that the observed strength of concept pair associations (Studies 1 and P1) from the priming literature predicted my capacity to observe (Studies 2, P2, and P5) or not (Studies 2, P3, P4, and P5) social priming effects. An attempt to strengthen a concept association and test the causal role of association strength (Study 3) resulted in no difference in concept associations and no priming. Overall, I find that social priming effects occur only for the most strongly related concept pairs and that easily measured explicit associations are useful for predicting these effects. I conclude that estimating the strength of underlying concept associations is critical for planning, interpreting, and reporting priming studies.

ACKNOWLEDGEMENTS

Completing a dissertation is hard but made much easier through the help of many people.

Thanks first to my committee members, who have made considerable contributions to my development as a researcher and whose comments have greatly improved this work. And especially to Dr. Kraus, Dr. Preston, and Dr. Cimpian who have been outstanding advisors and collaborators and who have supported my research.

This research could not have been done without the thousands of anonymous Mechanical Turk workers who participated in these studies. Their work accelerates the accumulation of psychological knowledge.

I am also grateful to the Graduate College at the University of Illinois for their support for this project through a Dissertation Completion Fellowship. And to the many colleagues who have read or listened to drafts of this work, especially the members of the Cognitive Development Lab and the Champaign Social Interaction Lab at the University of Illinois, whose feedback is directly reflected in this document.

Finally, I am thankful to those who provided social and emotional support during my doctoral studies. My husband, Erik Twombly, who agreed to move to a strange city in a strange state and wait patiently while I finished school. My non-academic friends who remind me that there is much more to the world, to life, and even to myself than regression, regression, regression. My academic friends who have inspired me to be better, shared in the necessity of forging forward through epistemological crisis, and helped me escape the everyday through hiking, cooking, synchronized vomiting, and adventurous travel. And my sincerest gratitude to my parents, who encouraged me to go off on this mad adventure.

TABLE OF CONTENTS

CHAPTER 1: INTRODUCTION

Priming—the idea that exposure to environmental stimuli activates related concepts that influence later cognitions, affect, and behaviors (e.g., Higgins, Rholes, & Jones, 1977)—has been a staple of social psychology for at least three decades. In one famous demonstration, after being incidentally exposed to words related to the elderly stereotype, people walked more slowly, acting out part of the stereotype's content (Bargh, Chen, & Burrows, 1996). Priming has been reported to occur across a wide variety of concept association pairs (e.g., POWER primes SELF-ESTEEM; GOD primes WATCHING, and TIME primes SELF-REFLECTION; Briñol, Petty, Valle, Rucker, & Becerra, 2007; Gervais & Norenzayan, 2012; Gino & Mogilner, 2014) and using a number of methods of prime exposure (e.g., recall of autobiographical events, subliminal exposure, sentence unscrambling; Baldwin & Holmes, 1987; Bargh, Bond, Lombardi, & Tota, 1986; Srull & Wyer, 1979) and response measurement (e.g., judgments, behavior, and reaction times; Arndt, Greenberg, Pyszczynski, & Solomon, 1997; Bargh et al., 1996, Fazio, Sanbonmatsu, Powell, & Kardes, 1986). The accumulated evidence has established the powerful and pervasive effect of priming on social behavior as a robust rule of social psychology. Its empirical basis is broad (covering all aspects of social psychology), deep (with careful consideration of mechanisms), and apparently sound (counting numerous conceptual replications of the most prominent phenomena).

Yet the recent history of social priming has been plagued by failures to replicate important effects (e.g., Earp, Everett, Madva, & Hamlin, 2014; Harris, Coburn, Rohrer, & Pashler, 2013; Johnson, Cheung, & Donnellan, 2014; Klein et al., 2014; Pashler, Coburn, & Harris, 2012) and criticism from cognitive psychologists (e.g., Doyen, Klein, Simons, & Cleeremans, 2014; Harris et al., 2013). Why was priming successful in original studies but not replications? Many explanations have been offered. Perhaps social priming itself is simply

implausible (Doyen et al., 2014). Or perhaps its effects are so context dependent that we should expect them to occur in only very specific circumstances (Cesario & Jonas, 2014). Given that priming has become such a fundamental phenomenon in social psychology, researchers are reluctant to abandon the phenomenon altogether (Molden, 2014). After all, even a dozen failures to replicate do not outnumber hundreds of successful priming studies. These controversies have prompted social psychologists to ask a fundamental yet unanswered question: How can social priming researchers predict the presence or absence of priming effects?

The answer to this question may stem from the nature of priming itself. Priming occurs as a result of spreading activation of related concepts in memory (Collins & Loftus, 1975): Perception of a stimulus in the environment activates a primed concept, which makes related concepts more accessible in later judgments (e.g., Meyer & Schvaneveldt, 1971). For example, perceiving words related to GOD activates related concepts, including GENEROSITY, which increases the tendency to behave prosocially (Shariff & Norenzayan, 2007). While social psychologists disagree about precisely how primes influence social processes—for instance, the active-self account links social priming effects to changes in self-concept (Wheeler, DeMarree, & Petty, 2007, 2014), whereas the situated inference model links them to mistaken attributions (Loersch & Payne, 2011, 2014)—they do agree that spreading activation is a fundamental mechanism of prime-to-behavior effects (e.g., Loersch & Payne, 2014; Wheeler et al., 2014). Thus, concept associations are the core of conceptual priming and may also be key to predicting the presence (or absence) of social priming effects and to interpreting the results of priming experiments.

1.1 A Brief Note on "Social" Priming

Priming as a psychological construct is used to describe a wide range of phenomena, all of which are supposedly united by the underlying mechanism of spreading activation. This

includes research in cognitive psychology finding that brief exposure to a word facilitates recognition of related words several milliseconds later (e.g., Neely, 1977) as well as research in social psychology finding that completing a form with a US flag background influences voting choices months later (Carter, Ferguson, & Hassin, 2011). These findings vary on several dimensions, the most prominent of which are the time scale of the effects (milliseconds to months) and the types of dependent variables (word recognition vs. reported intentions, beliefs, and behaviors). Recent controversies over priming have focused on only a subset of these phenomena characterized by timescales longer than milliseconds and by more complex social behavior. Such effects have been dubbed *social priming* (e.g., Harris et al., 2013, Kahneman, 2012) in contrast to *cognitive priming*, characterized by brief priming effects on cognitive dependent variables and which have enjoyed a history of early and positive replications (e.g., Segal, 1967). Throughout this text, I address issues related to predicting and interpreting social priming effects thus defined.

1.2 Concept Associations in Social Priming

The key argument of my proposal is that social priming, like cognitive priming (Fischler, 1977; Segal, 1967), should depend on and be predictable from concept association strengths. If, for example, undergraduate students at a given university do not associate GOD with WATCHING, activating the concept GOD in these students' minds should not activate thoughts related to WATCHING and, thus, they should not report greater self-consciousness or behave as if they are being watched any more than if they had not thought about GOD at all. In the absence of a sufficiently strong concept association, no priming effects should occur. And knowing the strength of a concept association in a given population should aid in the prediction of priming effects and the interpretation of priming studies.

Early work on automatic attitude activation suggests that this is true at least for attitude

priming: Priming an attitude object facilitates evaluation (i.e., positive or negative appraisal) only when people hold strong attitudes, or when they are trained to associate an attitude object with an evaluation (Fazio, 1993; Fazio et al., 1986; c.f., Bargh, Chaiken, Govender, & Pratto, 1992). In other words, the strength of the object-attitude link predicts the presence or absence of attitude priming effects. This research lends support to my central hypothesis, but it differs in important ways from work characterized as social priming by contemporary scholars. Namely, the dependent measures in these studies were (transformed) response latencies on the order of milliseconds. In other words, although these studies show that association strength influences priming of a social construct—attitude activation—they do not show that this dependency extends to the type of longer time scale social priming at the heart of current controversies.

In fact, little social priming research addresses concept associations explicitly. Surprisingly, priming researchers rarely assess or report information on concept associations. In one exception, individual differences in the degree that participants associated Black men with danger were assessed prior to examining whether these participants would respond to Black men with threat-related behaviors (Cesario, Plaks, Hagiwara, Navarrete, & Higgins, 2010; c.f., Bargh et al., 1996). This example is the exception rather than the rule: Most studies in social priming assume that the concepts are related in memory in the populations under study or use priming as a means of establishing that such an association exists (e.g., Gino & Mogilner, 2014; Mortensen, Ackerman, Becker, Neuberg, & Kenrick, 2010; Shariff & Norenzayan, 2007). Yet, social priming methods do not assess concept associations and so cannot establish them. This failure to assess concept associations complicates interpreting the results of social priming studies.

1.3 Interpreting Social Priming Results in Light of Concept Associations

Suppose a researcher has run a social priming study. She asked participants to think about death (Death Prime condition) or intense pain (Control condition) and then asked them to report

how much they felt in charge of their lives and their relationships. When she ran her analysis on the data, she found no difference between the groups on how in charge they felt. Why did she obtain this result? She considers two possibilities. First, she wonders if DEATH is simply unrelated to BEING IN CHARGE. Second, she considers that perhaps the materials were not quite right. The scale she used to measure feeling in charge might not be as reliable as she would like. Or the manipulation may have been too weak to sufficiently activate the concept of DEATH. Should she try to tweak the design and run the study again with a stronger manipulation and better measurement? Or should she cut her losses and move on to a new concept pair or a new project? Knowing whether DEATH is associated with BEING IN CHARGE could help her avoid running priming studies on unrelated concepts and arrive at clearer answers about the role of methods in failures to obtain effects. Compared to the resources involved in running a priming study, assessing concept associations is fast and easy and, when a concept association is weak, may save resources that would be spent tweaking a priming study one or more times before ultimately giving up. When a concept association is strong, researchers may be justified in considering that an initial study returning null results has done so for purely methodological reasons. And when it is weak, researchers can move on quickly and with little investment.

But what if the initial experiment had produced a significant effect? Would the researcher be justified in concluding that the manipulation changed feelings of being in charge through priming? No. There is no evidence that the result was obtained as a result of spreading activation from DEATH to BEING IN CHARGE. Such a result is consistent with a concept association based interpretation but does not substantiate it without either ruling out alternative accounts (e.g., that the task prompted participants to reflect on and learn about aspects of social relationships in a way that affected their feelings of being charge) or by providing evidence that

the concept association exists. Thus, measuring concept associations can save researchers resources that might otherwise be spent pursuing unlikely social priming effect and aid in the interpretation of studies that claim to find evidence for social priming.

Greater attention to concept associations in social priming can also aid in interpreting replication studies, especially when replications are conducted in new populations. Several recent replications of social priming effects have found no evidence for the original effect in their replication samples (e.g., Brandt, IJzerman, & Blanken, 2014; Doyen, Klein, Pichon, & Cleeremans, 2012; Earp et al., 2014; Fayard, Bassi, Bernstein, & Roberts, 2009; Harris et al., 2013; Johnson et al., 2014; Klein et al., 2014; Pashler et al., 2012), prompting researchers to search for explanations for the discrepancies between the original and replication findings. Do the replication results indicate that the original effect does not hold at all? Have the replication researchers discovered a previously unknown boundary condition on the effect? Were the materials appropriately adapted to the new population? Many answers to these questions have been proposed (e.g., Ferguson, Carter, & Hassin, 2014; Wentura & Rothermund, 2014; Williams, 2014), but there is little consensus regarding their accuracy. Had the original and replication researchers reported measures of the target associations in their populations, scholars could reconcile the opposing findings more easily. If the original researchers reported a strong association and the replication researchers a weaker one, this would support the possibility of a boundary condition: The effect may hold in the original population but not in the replication population due to some social moderator (e.g., a less religious replication population might have weaker associations for religious concepts). If, however, there were a strong concept association in both populations, the possibility of inappropriately adapted materials seems more likely. Thus, replications can best answer important questions about the stability and generalizability of priming effects when concept associations are measured.

1.4 Should Measuring Concept Associations Supplant Social Priming?

If psychologists use social priming to demonstrate concept associations, why conduct social priming studies at all when directly measuring concept associations does this better and more efficiently? Social priming studies provide real, important evidence about social phenomena—just not direct evidence about concept associations. Concept associations are part of a cognitive model of the structure of the mind. Representations stored in memory are "linked" by associations such that those concepts that share meaning or are frequently encountered together are "closer" together. Activation (e.g., through environmental stimuli) spreads more easily and quickly along these links when they are closer (Collins & Loftus, 1975). These associations are generally assessed through reaction time measures (e.g., Cattell, 1887; Greenwald, McGhee, & Schwartz, 1998; Segal, 1967) or by directly asking people about their concept associations (e.g., Nelson, McEvoy, & Schreiber, 2004; Rosch & Mervis, 1975). Social priming studies, in contrast, generally study the interrelationships among more complex behaviors—for example how completing a sentence unscrambling task influences the speed at which one walks (Bargh et al., 1996). These tasks do not assess or establish the presence of associations in the mind, but the theoretical account used to explain them relies on spreading activation among associated concepts.

Moreover, given a concept association, there is no reason necessarily to believe that it will have any specific influence on judgments or behaviors—or to know under what conditions that influence will occur or what its size will be. For example, that people associate DISPLAYING THE AMERICAN FLAG with REPUBLICANS (Carter et al., 2011) does not entail that exposure to the American flag will affect voting behavior two or more weeks later, nor does it establish how large such an effect should be. *This* is the purpose of social priming studies: to show whether, when, and how much spreading activation from incidentally encountered

concepts influences social judgments, decisions, and behavior. But to claim such influences are truly priming requires measuring concept associations.

1.5 The Present Research

Concept associations are the theoretical foundation of social priming and should therefore be central to priming research and theory. But how strong must concept associations be to observe a priming effect? This question has not yet been examined in the social priming literature. The present research explores the basis of priming in concept associations, focusing on the prediction that moderate to large concept associations are necessary for the primed concept to activate the target concept and influence subsequent judgments. This theory is tested in three studies.

First, I obtain estimates of the population-level strength of association for concept pairs from the social priming literature (Study 1) before testing whether priming occurs for a strong, moderate, and weak association (Study 2). These studies serve as an initial test of the idea that concept association strength should be predictive of the presence or absence of priming effects. If priming effects are related to concept association strength, priming should be observed for strong but not weak concept associations. The final study (Study 3) extends this work by attempting to experimentally manipulate the strength of a concept association prior to a priming task. This approach offers the benefit of causal inference about the role of concept association strength in priming effects. Five preliminary studies conducted as part of this project, providing evidence complementing Studies 1 and 2, are reported in Appendix A. Materials, anonymized data, and analysis code for all studies from this project are available from https://osf.io/6fm4a/. Full data, including demographic information, is available upon request from the author at ecsalomon@gmail.com.

CHAPTER 2: STUDY 1, MEASURING ASSOCIATIONS

Study 1 measured the strength of association for concept pairs from the social priming literature in a population commonly used for priming research: workers on Amazon's Mechanical Turk service. To estimate association strength, participants were asked to rate how strongly 15 concept pairs from the power priming literature were related. These association strengths will then be used to select of strong, moderate, and weak associations for later studies.

While some implicit attitude researchers have argued that people are unable to report their concept associations (e.g., Greenwald & Banaji, 1995), there is good evidence that they can. The Implicit Association Test (IAT) is designed to measure participants' implicit evaluative and concept associations (Greenwald et al., 1998)—exactly those associations they should be least able to report. Yet not only can people predict their scores on the IAT, they can do so beyond their guesses about how people in general will score (Hahn, Judd, Hirsh, & Blair, 2014). In other words, people have unique insight into their own associations. Moreover, early work on the availability heuristic found that people are highly accurate at estimating the cognitive accessibility of concepts within a category: The correlation between participants' guesses about performance on an availability task and actual performance was $r = .93$ (Tversky & Kahneman, 1973, Study 2). Given that people are accurate in reporting on the availability and associative strength of concepts, a self-report measure of concept associations, such as that used in Study 1, should provide a good estimate of actual associations.

All of the concept association pairs in Study 1 were selected from the power priming literature. Relying a single prime concept (POWER) allowed me to use a single manipulation in subsequent priming studies, increasing experimental control and allowing stronger inferences when comparing the effects of primes on strongly, moderately, and weakly associated concepts. Because the same experimental manipulation can be used to activate the primed concept for each

association, differences in the priming effects will not be due to differences in the effectiveness or strength of the priming manipulation.

2.1 Method

Participants. To provide precise estimates of concept association strength, I aimed for a final sample of at least 500 participants, excluding participants who take the survey from non-US IP addresses and/or fail an instructional manipulation check. Because power primes are common manipulations on Mechanical Turk (Meuller, Chandler, Paolacci, 2012), participation in Studies 1, 2, and 3 was restricted to workers who had completed 500 or fewer tasks on Mechanical Turk. To meet the target of 500 cases, I recruited 557 Mechanical Turk workers in exchange for $0.50. Of these, 12 were excluded for each of the reasons given above, leaving a final sample of 533 participants (281 women; 252 men). They averaged 33.52 years old (*SD* = 11.44), and most identified as European American (68%). Most participants (92%) had completed at least some college, and 48% reported an annual income of $35,000 or less. All studies run as part of this project used TurkGate software to prevent participants who had completed earlier studies in the same research line from participating in later studies.

Procedure. Participants completed a concept association measure followed by a demographics questionnaire.

The concept association measure asked participants to report how related 15 different concepts were to POWER using a slider to select any whole number from 0 (*not at all related*) to 100 (*strongly related*). Relatedness was defined as "how quickly or easily" POWER brings the target concept to mind. To reduce ambiguity about the concept of POWER (e.g., is it *having* power or *encountering* power?), participants were asked to "think about yourself having power. Imagine what it is like for you to have control and influence over others." Target concepts were presented in a random order for each participant. The 15 target concepts were:

- BEING IN CHARGE (Fast, Gruenfeld, Sivanathan, & Galinsky, 2009)
- SELF-ESTEEM (Briñol et al., 2007)
- TAKING ACTION (Galinsky, Gruenfeld, & Magee, 2003)
- APPROACHING OPPORTUNITIES (Smith & Bargh, 2008)
- PURSUING GOALS (Guinote, 2007)
- THINKING CAREFULLY (Smith, Dijksterhuis, & Wigboldus, 2008; Smith, Jostman, Galinsky, & van Dijk, 2008)
- BEING COMPETENT (Dubois, 2011)
- SEEING THINGS FROM OTHERS' PERSPECTIVES (Schmid Mast, Jonas, & Hall, 2009)
- OBJECTIFYING OTHERS (Gruenfeld, Inesi, Magee, & Galinsky, 2008)
- BEING PERSUASIVE (Lammers, Dubois, Rucker, & Galinsky, 2013)
- PRIDE (Schmid Mast et al., 2009)
- SELF-INTEREST (Chen, Lee-Chai, & Bargh, 2001)
- DISTRUSTING PEOPLE (Inesi, Greunfeld, & Galinsky, 2012)
- SOCIAL RESPONSIBILITY (Chen et al., 2001)
- SEX (Bargh, Raymond, Pryor, & Strack, 1995)
- STEREOTYPING PEOPLE (Lammers, Stoker, & Stapel, 2009)

The demographics questionnaire used in this and all subsequent studies asked participants to report several characteristics often thought to influence social cognition: gender, age, country of citizenship, religious affiliation, ethnicity, political orientation (a single 7-point item ranging from *very conservative* to *very liberal*), and subjective (i.e., ladder-based; Adler et al., 1994; Kraus, Piff, & Keltner, 2009) and objective (i.e., education and annual income; τ_B = -.36, z =

8.43, $p < .001$) socioeconomic status. An instructional manipulation check followed the measure of subjective social status. It presented the same ladder used to illustrate subjective social status in the previous question, but the question text was replaced with an instruction to "Please select '1 - Bottom rung' as your answer to this question." Participants were included in analyses if they followed this direction.

2.2 Results

Data for this project were analyzed in R version 3.2.3. Because several of the concept association distributions were skewed (skewness ranged from $g_1 = -2.25$ to $g_1 = 0.96$), medians were used as measures of central tendency and bias-corrected and accelerated bootstrapping (10,000 samples) was used to construct 95% equivalent confidence intervals. Using means rather than medians does not substantially change the interpretation of results, though it does restrict the range of central tendencies (see Appendix B for a figure comparing median and mean concept association strengths).

Figure 2.1 (page 14) depicts the median concept association scores for the 15 target concepts from the power priming literature included in this study. Mechanical Turk workers' reported strength of associations varied greatly, ranging from 15 (POWER and STEREOTYPING PEOPLE) to 97 (POWER and BEING IN CHARGE).

2.3 Conclusions

Concept associations studied in power priming research vary greatly in their strength, suggesting that power priming research may be investigating some constructs only weakly associated with power and that failures to replicate such studies on Mechanical Turk might be expected. These results are consistent with results from a preliminary study, Study P1 (see Appendix A), which found that associations for concept pairs from the social priming literature also vary greatly, covering essentially the full range of the scale. Thus, power priming research

seems to be a reasonable representative of social priming for the purpose of studying concept associations. Studies 2 and 3 investigate the hypotheses that (1) more strongly associated concept pairs will be better candidates for priming effects than will more weakly associated concept pairs and (2) increasing a concept association will result in a stronger priming effect.

2.4 Figure

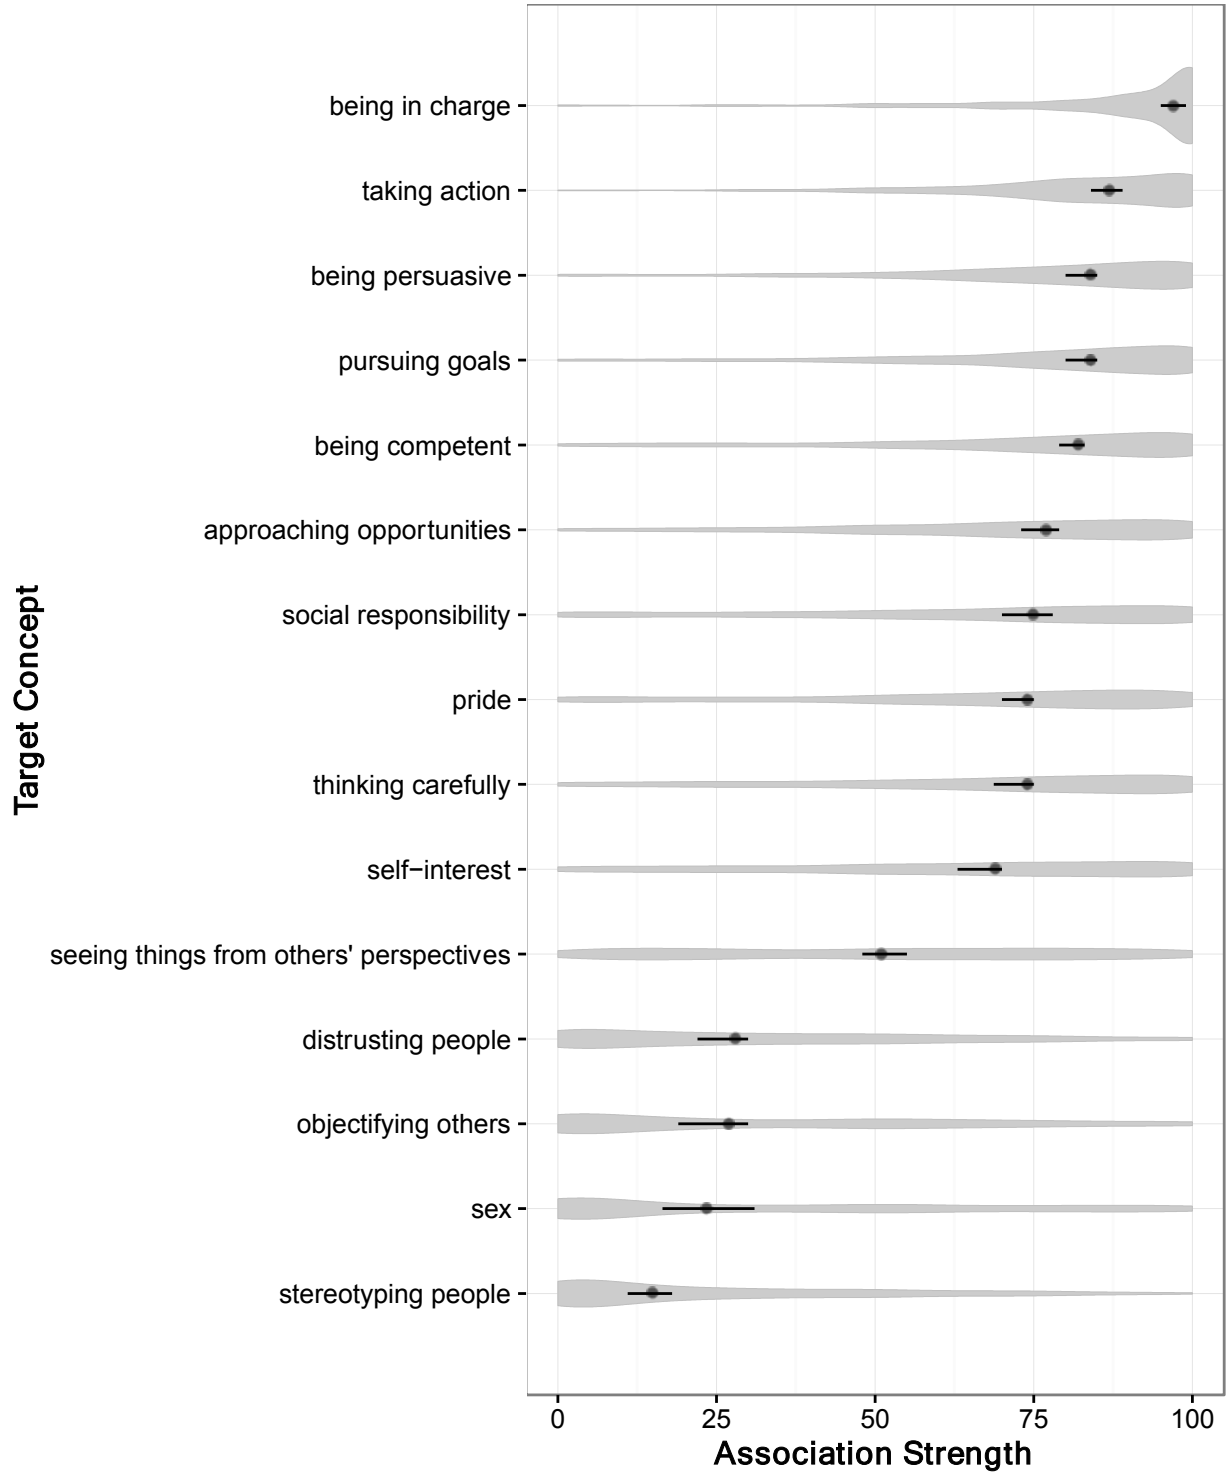

Figure 2.1. Median association ratings (circles) for 15 concept pairs from the power priming literature, Study 1. Error bars represent 95% bias-corrected and accelerated bootstrapped confidence intervals based on 10,000 samples. Gray regions are violin plots showing the kernel smoothed density for each distribution.

CHAPTER 3: STUDY 2, PRIMING ASSOCIATIONS OF DIFFERENT STRENGTHS

Study 2 tested the prediction that a strong concept association would yield a priming effect but a weak one would not. It also explored the priming potential of a moderate concept association. Participants were primed with high or low power and were asked to make judgments relating to a strongly (TAKING ACTION; *median association in Study 1* = 87), moderately (SEEING THINGS FROM OTHERS' PERSPECTIVES; *median* = 51), and weakly associated concept (DISTRUSTING PEOPLE; *median* = 28). If social priming is predictable from concept associations, I should observe an effect of power priming on judgments related to TAKING ACTION but no effect on judgments related to DISTRUSTING PEOPLE.

Moreover, measuring differences on the strongly, moderately, and weakly related concepts in a single study allows me to test competing models of the relationship between association strength and social priming effects. According to a linear model (e.g., Coney, 2002), the strength of a priming effect increases linearly with association strength, such that I should observe the largest difference for TAKING ACTION and the smallest difference for DISTRUSTING PEOPLE, with SEEING THINGS FROM OTHERS' PERSPECTIVES falling between those two. A threshold model, in contrast, predicts that social priming effects should occur only when associations exceed some strength threshold. These models differ primarily in the predictions they make about moderately strong concept associations. Under the linear model, such associations should produce priming effects (though smaller than those for strong associations). Under a threshold model, predictions about priming effects for moderate concept associations depend on the level of the threshold. If the threshold is low, such effects should approximately the same size as strong concept associations, but if the threshold is high, no priming effects should be observed for moderately associated concepts.

3.1 Method

Participants. I again aimed to collect data from at least 500 participants after exclusions for non-US IP addresses and failing an instructional manipulation check. This provides 80% power to detect an effect size of at least $d = 0.29$ at $\alpha = .017$ (Bonferroni-corrected for three comparisons), which is slightly smaller than the average effect of behavioral priming (Weingarten, Chen, McAdams, Yi, Hepler, & Albarracín, 2016). To meet this goal, I recruited 626 Mechanical Turk workers, who were paid $2.00 each for completing the survey. 19 participants were excluded for completing the survey from non-US IP addresses and 29 were excluded for failing an instructional manipulation check (including one participant excluded for both reasons), leaving 579 participants (357 women; 222 men) included in analyses. They averaged 32.06 years old ($SD = 13.17$). Most identified as European American (68%). Most participants (89%) had completed at least some college, and 51% reported an annual income of $35,000 or less.

Procedure. Participants completed a power priming task; dependent measures designed to test their desire to take action, their ability and motivation to take others' perspectives, and their trust for others; and the same demographics questionnaire used in Study 1.

Priming task. Participants were randomly assigned to one of two power priming conditions. They recalled a situation in which they had power over someone else (High Power) or in which someone else had power over them (Low Power; Galinsky et al., 2003). They were guided to vividly relive this experience during a 60 second visualization task (e.g., ". . .try to zoom in and get a close-up, focused impression. Hold this image for a little while. Imagine talking with the person. Try to feel them there with you") and then describe the memory and how it made them feel (Baldwin & Holmes, 1987). To ensure that participants did not simply copy and paste pre-written power essays, JavaScript was used to disable pasting on the priming task.

Dependent measures. Seven questions assessed participants' action and inaction goals (e.g., "If I could, I would go work out after this study"; "If I could, I would gladly take a nap after this study") on scales ranging from 1 = *not at all* to 7 = *extremely* (Albarracín, 2009).

Participants' ability and desire to take others' perspectives was measured with two scales. Seven items from the perspective-taking subscale of the Interpersonal Reactivity Index (IRI) asked participants to rate how much they try to see things from others' perspectives (e.g., "I find it difficult to see things from the 'other guy's' point of view") on scales ranging from 1 = *does not describe me at all* to 5 = *describes me very well* (Davis, 1980). Several items were slightly reworded to capture state rather than trait variability. For example, "I try to look at everybody's side of a disagreement before I make a decision," was reworded as, "I'd want to look at everybody's side of a disagreement before I make a decision." In addition, 10 items from the Reading the Mind in the Eyes Test Revised Version asked participants to guess which of four emotions was expressed in a photograph of a person's eyes (Baron-Cohen, Wheelwright, Hill, Raste, & Plumb, 2001). Each item included one correct and three incorrect answers.

Participants also completed two measures of their (dis)trust for others. Six items from the World Values Survey (WVS) asked participants how much they trust people from different groups (e.g., "Your family"; "People you meet for the first time") on scales ranging from 1 = *trust completely* to 4 = *do not trust at all* (World Values Survey Association, 2012). In addition, participants played a single round of the Trust Game, an economic game designed to measure trust (Berg, Dickhaut, & McCabe, 1995). The general structure of this game is as follows:

1. One player (A) is given some money (or other item of value) and may share it with another player (B);
2. All money given to B is tripled by the experimenter;
3. B returns as much of the tripled money to A as s/he wishes, and the game ends.

All participants in Study 2 were assigned to the role of A and given 1000 tickets to a raffle for a $50 bonus payment. Participants were given the rules to the game and told that they were paired with another survey-taker in the role of B—though there was, in reality, no such other player. Trust was measured by how many tickets they entrusted to (the non-existent) B. After the survey, all participants were given equal entry to the raffle, which was awarded to a randomly selected participant.

The three sets of dependent measures (action goals, perspective-taking, and trust) were presented in a random order. The relative orders of the two perspective-taking measures and of the two trust measures were also randomized, as were the orders of items within each measure.

3.2 Results

The action/inaction goals scale was scored by keying all items so that higher values indicate greater action or weaker inaction goals and then taking their average ($\alpha = .84$). The IRI was similarly scored by keying all items so that higher scores indicate greater intentions to take others' perspectives and then averaging ($\alpha = .83$). The Reading the Mind in Eyes Test was scored as the proportion of correct responses. The WVS trust measure was scored by keying items so that greater numbers indicated greater distrust and averaging them ($\alpha = .74$). The Trust Game was scored as the number of tickets kept by the participant so that higher numbers indicate greater distrust. Table 3.1 (page 22) shows correlations among the dependent measures. Correlations between the two perspective-taking measures, $r(577) = .00$, $p = .923$, and between the two trust measures, $r(577) = .09$, $p = .024$, were weak, so these measures were treated separately in the remaining analyses.

If strong but not weak concept associations produce social priming effects, we should expect to see a significant difference between conditions for the action/inaction goals measure but not for either of the trust measures. Likewise, if priming occurs for moderately related

concept pairs, we would expect to see significant differences between conditions on the perspective-taking measures. Means by condition, *t*-tests, and Cohen's *d*s for all dependent measures by condition are reported in Table 3.2 (page 22). All *t*-tests use Welch's correction for unequal variances. To control errors via Bonferroni correction, α is set to .01. Consistent with my prediction, participants in the High Power condition reported significantly stronger action goals than did participants in the Low Power condition, $t(575.31) = 3.13$, $p = .002$, $d = 0.26$ [0.09, 0.42]. There were no differences between conditions on the perspective-taking or trust dependent variables. Most of the effect size estimates for these comparisons are close to 0. However, the trust game shows a small, non-significant difference such that participants in the High Power condition showed entrusted fewer tickets to their partners than participants in the Low Power condition, consistent with an association between POWER and DISTRUSTING PEOPLE but contrasting results from Study P5 in Appendix A.

I formally tested the prediction that higher association strengths lead to larger priming effects with three exploratory linear mixed effects models. In these models, the critical test of my hypothesis is the fixed effect for the interaction between Association and Priming condition: Does priming emerge as a function of association strength? However, the models represent different predictions about the way association strength relates to priming effects.

In the first model, I tested the prediction that priming effects increase linearly with association strength—the more associated concepts are, the more priming will occur. The five dependent measures were assigned a centered and scaled contrast code based on the median observed association strengths in Study 1: -0.21 = POWER and DISTRUSTING PEOPLE (used for the WVS trust measure and the Trust Game), 0.02 = POWER and SEEING THINGS FROM OTHERS' PERSPECTIVES (used for the IRI perspective-taking measure and the Reading the Mind in the Eyes Test), and 0.38 = POWER and TAKING ACTION (used for the

Action/Inaction Goals scale). A change from -0.50 to 0.50 on this scale is equivalent to the difference between an association of 0 and an association of 100 on the scale used in Study 1. To create a common scale for the dependent measures, all five were standardized prior to analysis. If priming effects increase linearly with association strength, I should observe a positive interaction, indicating that the standardized difference between the Low Power and High Power conditions increases as association strength increases. As shown in Table 3.3 (page 23), this fixed effect was indeed positive but did not reach significance.

In the second and third models, I tested the prediction that priming effects occur only for associations stronger than some threshold. Each of these set the threshold in a different place. The second model tests the prediction that social priming effects occur for all but the weakest concept associations. In this model, the threshold was set between the weak and moderate concept association. The dependent variables were coded as 0.40 (action goals and perspective-taking measures) or -0.60 (distrust measures). As shown in Table 3.4 (page 23), the fixed effect for the Association by Priming condition interaction was approximately 0 in this model, suggesting that a low threshold for association strength does not represent these data well.

The third model tests the prediction that social priming effects emerge only for strong concept associations. In this model, the dependent variables were coded as 0.80 (Action/Inaction goals scale) or -0.20 (all others). As shown in Table 3.5 (page 23), the fixed effect for the Association by Priming condition interaction was positive and significant at $p < .05$ in this model, suggesting that concept associations may need to exceed a high strength threshold for social priming effects to emerge. However, applying a Bonferroni correction to the three tests of the Association x Priming condition interaction ($\alpha = .017$) reduces this interaction to non-significance.

3.3 Discussion

Study 2 tested the hypothesis that strongly, but not weakly, associated concepts would produce priming effects. Support for this hypothesis was found in *t*-tests, which showed a significant difference only for the strongest association. Participants primed with high power showed greater action goals compared to those primed with low power. No significant priming effects were observed for moderate (POWER and SEEING THINGS FROM OTHERS' PERSPECTIVES) or weak (POWER and DISTRUSTING PEOPLE) concept associations. These results are consistent with those from preliminary research conducted as part of this project (see Appendix A for details). In Studies P2 and P3, priming was observed for a strongly (POWER and BEING IN CHARGE, Study P2) but not a weakly (TIME and SELF-REFLECTION, Study P3) associated concept pair, and in Study P4, there was little evidence for social priming for a moderately associated concept pair (GOD and WATCHING). Study P5 used a method very similar to that of Study 2 but with a different set of concept pairs. Its results parallel those presented here, though the difference for the strong concept pair did not reach significance.

In models testing different predictions about the relationship of association strengths to priming effects, I found some evidence for a threshold model: Priming effects appear to emerge only once the underlying concept association reaches a high critical level of strength. This result parallels that in Study P5 as well.

Overall, these studies suggest that explicit measures of concept associations can be used to predict the presence and absence of priming effects—and that only the strongest concept associations are good candidates for social priming.

3.4 Tables

Table 3.1

Correlations Among Dependent Measures, Study 2

Variable	*1*	*2*	*3*	*4*
1. Action/Inaction Goals				
2. IRI Perspective-Taking	.16***			
3. Reading the Mind in the Eyes	-.03	.00		
4. WVS Trust	.08*	.06	.07	
5. Trust Game	.03	.15***	-.01	.09*

Note. N = 579. IRI = Interpersonal Reactivity Index. WVS = World Values Survey.
* $p < .05$. *** $p < .001$

Table 3.2

Means (Standard Deviations), t-*Tests, and Cohen's* d*s for Dependent Variables, Study 2*

	Low Power		*High Power*						95% *CI* for *d*	
Variable	(*n* = 299)		(*n* = 280)		*t*	*df*	*p*	*d*	*LL*	*UL*
Action/Inaction Goals	4.11	(1.30)	4.43	(1.15)	3.13	575.31	.002	0.26	0.09	0.42
IRI Perspective-Taking	3.82	(0.59)	3.83	(0.65)	0.15	562.25	.878	0.01	-0.15	0.18
Reading the Mind in the Eyes	0.76	(0.15)	0.76	(0.14)	-0.23	577.00	.817	-0.02	-0.18	0.14
WVS Trust	2.09	(0.43)	2.09	(0.41)	0.15	576.82	.882	0.01	-0.15	0.18
Trust Game	512.80	(300.14)	566.13	(309.28)	2.10	571.76	.036	0.18	0.01	0.34

Note. N = 579. IRI = Interpersonal Reactivity Index. WVS = World Values Survey. Action/Inaction Goals are scored on a 1 to 7 scale. IRI Perspective-Taking is scored on a 1 to 5 scale. Reading the Mind in the Eyes test is scored as the proportion correct (out of 10). WVS Trust is scored on a scale from 1 (most trust) to 4 (least trust). Trust Game is scored as the number of tickets (out of 1000) kept by the participant. *T*-tests use Welch's correction for unequal variances. To correct for multiple comparisons, α is set to .01.

Table 3.3

Linear Mixed Effects Model Testing the Prediction of Linear Increase in Social Priming Effects with Concept Association Strength.

		95% *CI* for *b*			
Predictor	*b*	*LL*	*UL*	*t*	*p*
Priming Condition	0.09	0.01	0.16	2.33	.020
Linear Association Contrast	0.00	-0.17	0.17	0.00	1.000
Interaction	0.26	-0.08	0.60	1.48	.140

Note. $N = 579$. *P*-values are based on the Satterthwaite approximation for degrees of freedom. Confidence intervals are based on profile likelihood approximation.

Table 3.4

Linear Mixed Effects Model Testing the Prediction that Social Priming Effects Occur When Concept Associations Exceed a Low Strength Threshold.

		95% *CI* for *b*			
Predictor	*b*	*LL*	*UL*	*t*	*p*
Priming Condition	0.09	0.02	0.16	2.36	.019
Low Association Threshold Contrast	0.00	-0.08	0.08	0.00	1.000
Interaction	-0.01	-0.16	0.14	-0.09	.927

Note. $N = 579$. *P*-values are based on the Satterthwaite approximation for degrees of freedom. Confidence intervals are based on profile likelihood approximation.

Table 3.5

Linear Mixed Effects Model Testing the Prediction that Social Priming Effects Occur When Concept Associations Exceed a High Strength Threshold.

		95% *CI* for *b*			
Predictor	*b*	*LL*	*UL*	*t*	*p*
Priming Condition	0.09	0.01	0.16	2.32	.021
High Association Threshold Contrast	0.00	-0.09	0.09	0.00	1.000
Interaction	0.21	0.03	0.39	2.28	.023

Note. $N = 579$. *P*-values are based on the Satterthwaite approximation for degrees of freedom. Confidence intervals are based on profile likelihood approximation.

CHAPTER 4: STUDY 3, MANIPULATING A CONCEPT ASSOCIATION

Study 3 was designed to test the causal prediction that strong concept associations drive priming effects. To do this, I attempted to strengthen the weak association between POWER and DISTRUSTING PEOPLE using a novel concept association training program modeled after the conditioning paradigms of cognitive psychology (e.g., Pavlov, 1927/2003). The training task was designed to increase the behavioral pairing of having power over another agent and being more distrustful of others. I expected that the training task would successfully pair situational power with distrustful behavior, which I predicted would increase (1) the corresponding concept association and (2) the subsequent priming effect of having power on distrust.

To test this prediction, I designed a concept association training task wherein participants completed a modified multi-round trust game in which their partners' behaviors depended on (or were independent of) the amount of power they held on a given round. I predicted that those participants for whom having power was paired with untrustworthy partners would develop a stronger association between POWER and DISTRUSTING PEOPLE. After the training task, participants completed the power priming manipulation and trust measures from Study 2. If concept association strength is causally related to social priming, the priming manipulation should affect trust only for participants who strongly hold the POWER- DISTRUSTING PEOPLE association. This study was pre-registered on the Open Science Framework (https://osf.io/nm39c/).[1]

4.1 Method

Participants. I aimed to collect 400 participants for this study. This sample size provides 80% power to detect a difference of $d = .28$ in concept associations between participants trained

[1] The pre-registered analysis code contained an error in assigning condition codes to long format data. This was corrected before analyses were run.

to associate POWER with DISTRUSTING PEOPLE and those not trained and 80% power to detect an effect of $f^2 = 0.03$ in a regression model predicting trust from training condition, priming condition, and their interaction. To ensure that I collected data from at least 400 participants who met pre-registered inclusion criteria (i.e., who took the study from a US IP address, passed an instructional manipulation check, and completed the training task), I recruited 516 participants from Mechanical who were paid $2.50 each for participation.

17 participants were excluded for completing the survey from non-US IP addresses, 10 were excluded for failing an instructional manipulation check, and 60 were excluded for not completing the training task (including three participants excluded for multiple reasons), leaving 432 participants (243 women; 185 men; 4 other genders) included in analyses. They averaged 31.95 years old ($SD = 9.92$). Most identified as European American (72%). Most participants (91%) had completed at least some college, and 49% of participants reported an annual income of $35,000 or less.

Procedure. Participants completed a concept association training task, the power priming task from Study 2, the trust measures from Study 2, and a demographics questionnaire.[2]

Training task. Participants first completed a training task that consisted of a 30-round modified trust game. They were again told that they were playing for tickets to a $50 raffle. Figure 4.1 (page 32) illustrates gameplay for the training task. Participants were always assigned to the role of A (giver) and given 100 tickets to play with on each round. On each round, participants were purportedly matched with a randomly selected participant who had been assigned the role of B; however, B was always a simulated computer partner.

[2] This questionnaire differed from the one used in Studies 1 and 2 in that it allowed participants to write open ended responses for their gender identities, rather than selecting from only two options (Male or Female). Otherwise, the questions and response options were the same. Gender identities were coded into three categories: (cis- or trans-) Women, (cis- or trans-) Men, and Non-binary (e.g., genderfluid, agendered).

To manipulate power within the game, A and B took turns as Controller. The Controller had the option to punish the other player after their decision by eliminating one third of the tickets they kept (*Z* in Figure 4.1). Participants were Controller on half of the rounds (high participant power), and on the other half, B was Controller (low participant power). Controller assignments were blocked and the blocks randomly ordered so that participants completed 15 rounds as Controller followed by 15 rounds with their partners as Controller or vice versa.

To manipulate the association between POWER and DISTRUSTING PEOPLE, participants were randomly assigned to one of two Training conditions, which differed in the rules governing the number of tickets returned by B (*Y* in Figure 4.1). In the Association condition, the trustworthiness of partners varied as a function of Controller assignment. When participants in the Association condition were Controller, their partners behaved stingily, returning between 0% and 10% of the tickets they received. When the same participants were at the mercy of a Controller partner, their partners behaved generously, returning between 60% and 70% of the tickets they received. In the No Association, partners played fairly, returning between 45% and 55% of the tickets received, no matter who was Controller.

A demo version of the training manipulation, consisting of four rounds can be played at http://erikasalomon.com/trust/.

Concept association measure. To assess whether the training task successfully manipulated the association between POWER and DISTRUSTING PEOPLE, participants completed the same concept association measure from Study 1, which included the focal concept association as well as several others to disguise the association of interest.

Priming task. Participants were randomly to the same High Power or Low Power priming tasks used in Study 2.

Trust measures. As in Study 2, participants completed the World Values Survey trust

measure and an unmodified trust game with 1000 raffle tickets.

4.2 Results

Manipulation checks. Two methods were used to assess the manipulation of partner trustworthiness in the game.

First, if partners in the Association but not in the No Association condition were thought by participants to behave in differentially trustworthy ways based on Controller assignment, this should be reflected in interactions between Controller assignment and Training condition on the number of tickets given to partners. Participants in the Association condition—but not those in the No Association condition—should give fewer tickets (show less trust) when they are Controller than when their partners are Controller. To test this, I fit four nested linear mixed effect models. The base model included crossed random effects for participants and Controller assignments with time (game round) nested within each of these.[3] Model fit was improved by adding fixed effects for Controller assignment (-0.5 = Partner, 0.5 = Participant; $\chi^2(1) = 19.38$, $p < .001$), Training condition (-0.47 = No Association, 0.53 = Association; $\chi^2(1) = 30.43$, $p < .001$),[4] and their interaction ($\chi^2(1) = 31.28$, $p < .001$). The final model, including the interaction term, is summarized Table 4.1 (page 36). Figure 4.2 (page 33) illustrates the expected interaction pattern, with differential trust in the Association condition but not in the No Association condition.

Second, I used growth models (specified as exploratory in the pre-registration) to examine whether participants' responses reflected learning over time. If participants in the Association condition were learning the different patterns of partner responses across rounds,

[3] The model specified in the pre-registration used mean level, rather than round level, outcome data. The fixed effect of the Controller assignment by Training condition in this model was significant using the Satterthwaite approximation for degrees of freedom, $b = 19.12$ [15.11, 23.14], $t = 9.34$, $p < .001$.

[4] The weighted effect codes used for Training condition reflect slight asymmetry in the random assignment of participants to conditions. This is not an issue for the within-subjects factor Controller assignment. See Table 4.3 (page 37) for cell sizes for between participant conditions.

their behavior should change as a function of how many rounds they have completed, such that the number of tickets they give should increase over time when their partners are Controller (increasing trust to generous partners) and decrease over time when they are Controller (decreasing trust to stingy partners). Participants in the No Association should not show this differential learning pattern as their partners displayed fair behavior regardless of Controller assignment. To test this prediction, I added new random and fixed effects to the previous model to build nine new nested models. I first added a variance component for time^2 ($\chi^2(1) = 758.37$, $p < .001$). Model fit was further improved by adding fixed effects for time ($\chi^2(1) = 25.91$, $p < .001$); time^2 ($\chi^2(1) = 23.99$, $p < .001$); the time by Controller assignment interaction ($\chi^2(1) = 46.26$, $p < .001$); the time by Training condition interaction ($\chi^2(1) = 15.12$, $p < .001$); the three-way interaction among time, Controller assignment, and Training condition ($\chi^2(1) = 87.28$, $p < .001$); the time^2 by Controller assignment interaction ($\chi^2(1) = 25.14$, $p = .001$); the time^2 by Training condition interaction ($\chi^2(1) = 10.24$, $p < .001$); and the three-way interaction among time^2, Controller assignment, and Training condition ($\chi^2(1) = 53.25$, $p < .001$). Fixed effects coefficients for the final model, including both three-way interactions, are reported in Table 4.2 (page 36). As illustrated in Figure 4.3 (page 34), this model fits the observed data well and shows the expected learning curves. Participants in the Association condition decreased in trust over time when they were the Controller but increased when their partners were Controller. Participants in the No Association condition show flatter learning curves that differ little between Controller assignments.

Overall, these results show that the training task in the Association condition successfully manipulated the trustworthy behavior of partners as a function of power (Controller assignment).

Changes in concept associations. To test the hypothesis that the training task would induce changes in concept associations, I examined whether participants in the Association

condition reported a higher POWER- DISTRUSTING PEOPLE association than participants in the No Association condition. I computed a bias-corrected and accelerated bootstrapped 95% equivalent confidence interval (10,000 samples) for the difference in medians, which showed no evidence for a difference, *difference* = 1 [-10, 7]. As Figure 4.4 (page 35) illustrates, participants in both Training conditions reported weak associations between POWER and DISTRUSTING PEOPLE.

Priming. The WVS trust measure (α = .74) and the Trust Game were scored as in Study 2 so that higher numbers indicate greater distrust. They again correlated weakly, $r(430) = .14$, $p = .003$, and were treated separately in the remaining analyses with α set to .025. Table 4.3 (page 37) displays observed means for these variables.

The effect of concept associations on priming was examined in two ways. First, as specified in the pre-registration, I fit linear models predicting each dependent variable from Training condition (-0.47 = No Association, 0.53 = Association), Priming condition (-0.49 = Low Power, 0.51 = High Power), and their interaction. If participants in the Association condition developed a stronger association between POWER and DISTRUSTING PEOPLE relative to those in the No Association condition, Priming condition should have a larger effect on their trust. However, given that participants in the two Training conditions did not differ in their ratings of this association, it would be surprising to find this interaction. Consistent with this expectation, no interaction was observed for World Values Survey trust scores or number of tickets kept in the unmodified Trust Game (Table 4.4, page 38). Notably, the estimated interactions for the two trust measures are in *opposite* directions, with the interaction for the World Values Survey trust measure in the predicted direction. The interaction for the Trust Game, however, suggests that participants in the Association condition showed a *more negative* difference in distrust between the High and Low Power conditions. This is confirmed by the

means in Table 4.3, which show that, in the Association condition, participants who recalled a High Power experience kept *fewer* tickets for themselves than did participants who recalled a Low Power experience relative to participants in the No Association condition. This is especially surprising given that participants in the Association condition showed strong differentiation in their trusting behavior based on Controller (power) assignment in the training task, which was based on the Trust Game.

Second, given the failure of the Training manipulation to change concept associations, I ran two linear models, not specified in the preregistration, predicting each of the dependent variables from participants' reported concept associations, Priming condition, and their interaction. If participants who hold this association strongly exhibit a priming effect but those who hold it weakly do not, this should be reflected in the interaction term. However, as illustrated in Figure 4.4, few participants reported strong associations, weakening the ability to observe a priming effect in this range, and no interaction was observed for either dependent variable (Table 4.5, page 38). Notably, however, the interactions from these models were in the same directions (though smaller in magnitude) as those in the models reported in Table 4.4, with the interaction for the World Value Survey trust measure in the expected direction and for the Trust Game in the opposite direction to the prediction.

4.3 Discussion

Study 3 aimed to test the causal relationship between concept associations and social priming effects. I attempted to manipulate a concept association by training participants to expect others to be untrustworthy when they have power and trustworthy when they do not. Although the manipulation was successful in changing participants' trusting or distrusting behavior based on power dynamics, neither of the hypotheses were confirmed. Participants' associations between POWER and DISTRUSTING PEOPLE were unchanged, and the manipulation did not

significantly affect priming of distrust from a power manipulation. Moreover, measured concept associations did not predict presence, absence, or size of priming effects.

4.4 Figures

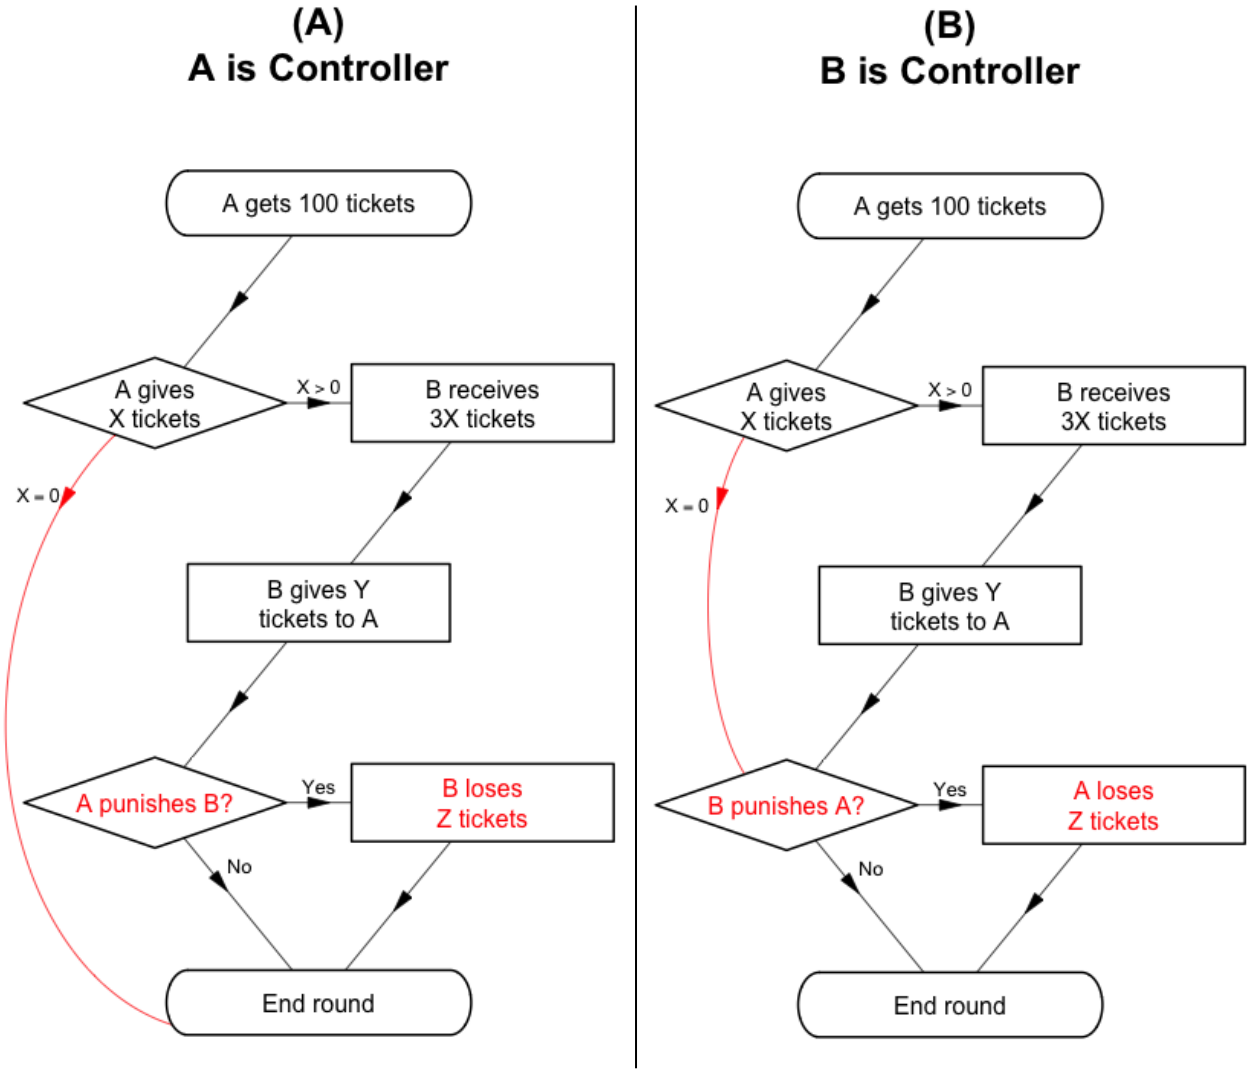

Figure 4.1. Flowcharts depicting modified Trust game rules when (A) A is Controller and (B) B is Controller. Red color indicates differences between (A) and (B). Participants are always assigned the role of A. B is always simulated. Values of *Y* depend on participants' assignments to Training condition (see text). *Z* is equal to one-third of *X* in (A) and one-third of *Y* in (B).

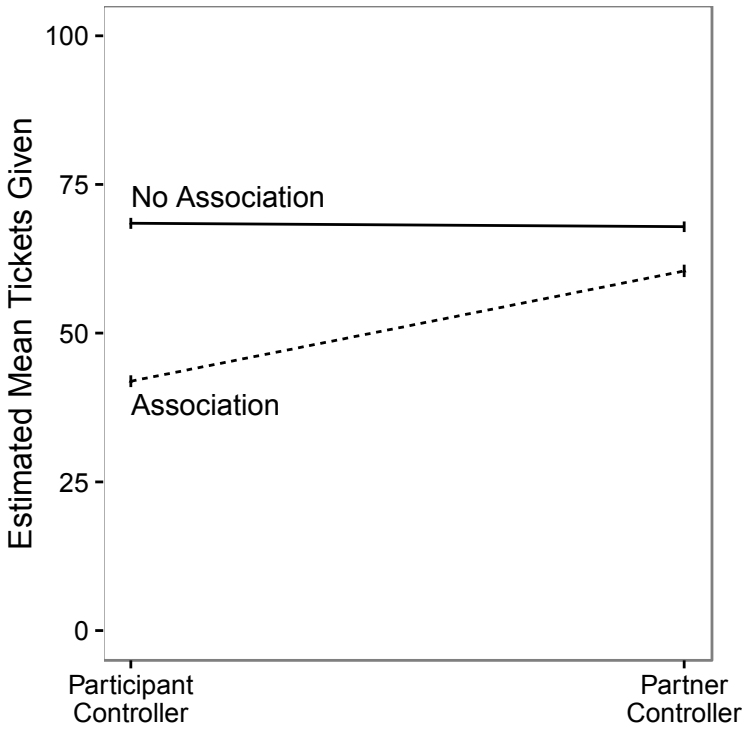

Figure 4.2. Estimated mean number of raffle tickets entrusted to partner across 15 rounds of a modified trust game as a function of power (Controller assignment), Study 3. Whereas participants in the No Association condition show equal levels of trust regardless of power, participants in the Association condition showed less trust when they had power than when their partners had power. Error bars represent 95% confidence intervals.

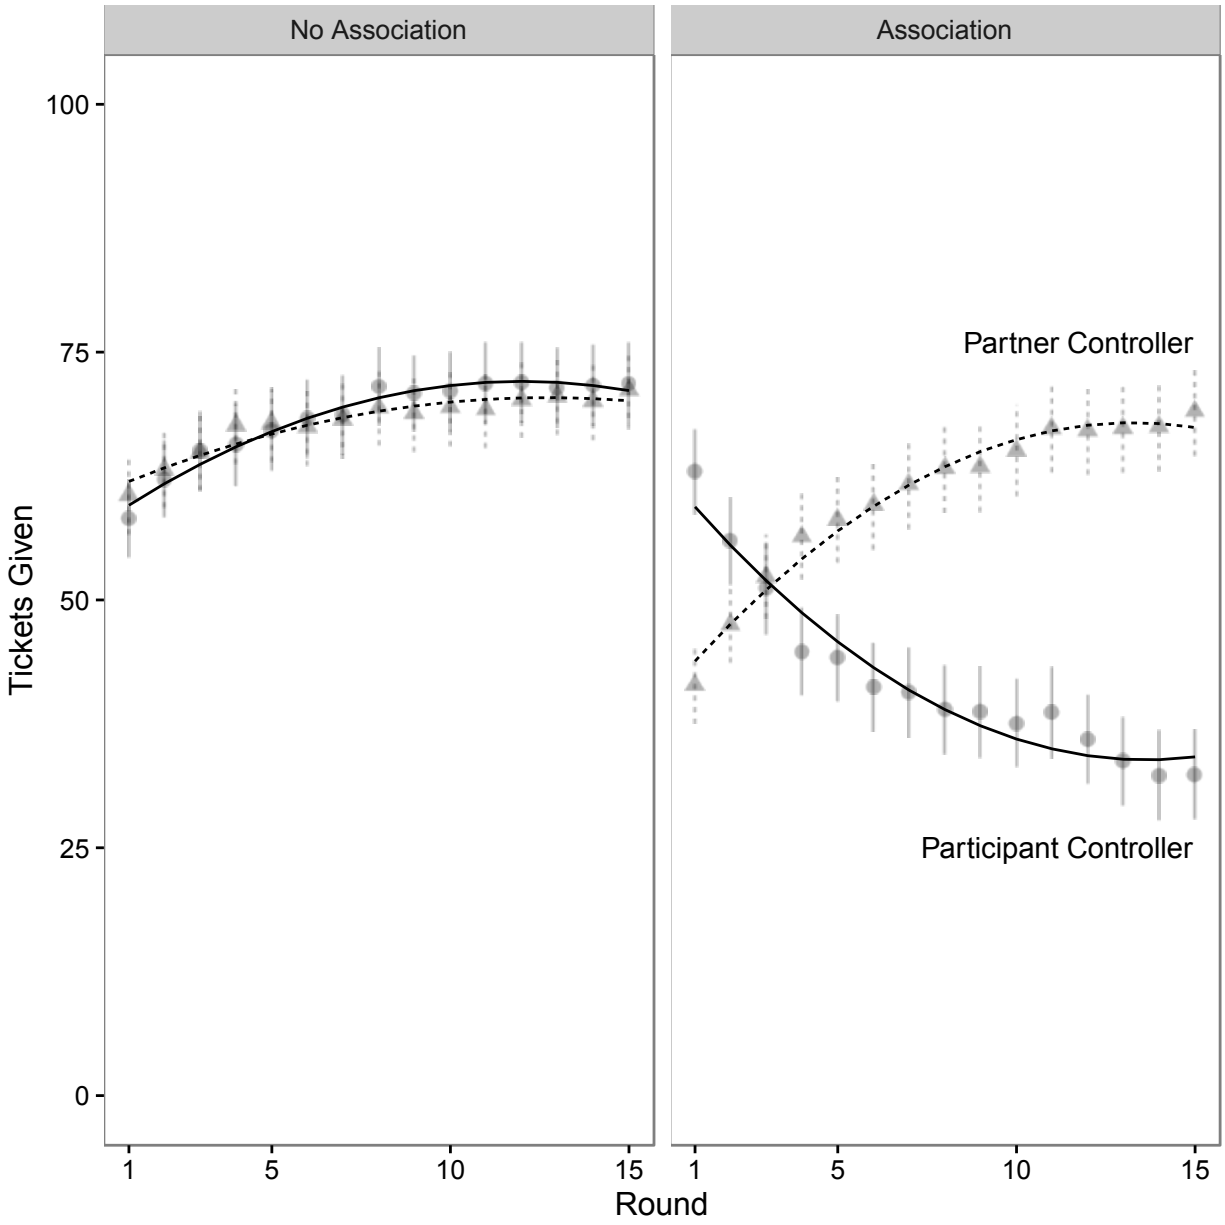

Figure 4.3. Fitted growth curves (lines) and observed values (points) with 95% confidence intervals showing learning over time in modified trust game with punishment, Study 3. Participants in the Association condition—but not the No Association condition—experienced differential partner behavior based on which player had power (Controller assignment).

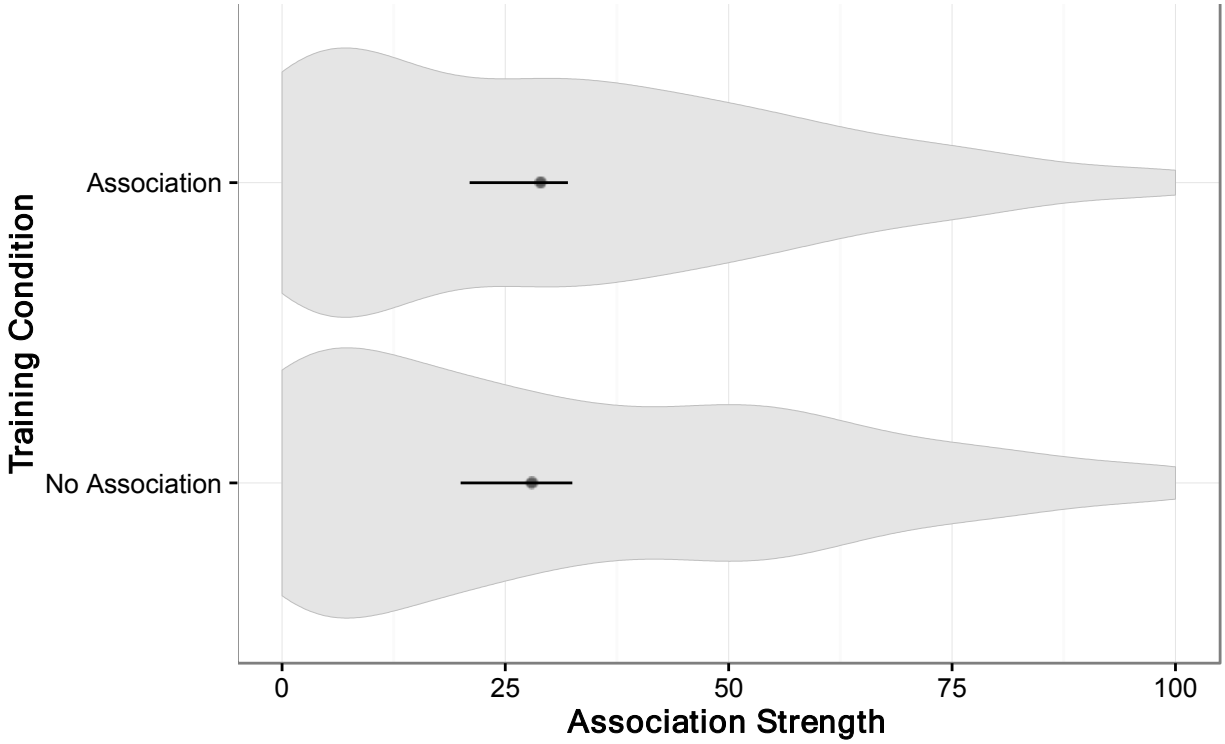

Figure 4.4. Median association ratings (circles) for POWER and DISTRUSTING PEOPLE, Study 3. Error bars represent 95% bias-corrected and accelerated bootstrapped confidence intervals based on 10,000 samples. Gray regions are violin plots showing the kernel smoothed density for each distribution.

4.5 Tables

Table 4.1

Linear Mixed Effects Model Predicting Mean Number of Tickets Given from Controller Assignment, Training Condition, and Their Interaction.

		95% *CI* for *b*			
Predictor	*b*	*LL*	*UL*	*t*	*p*
Intercept	58.47	55.85	61.08	50.55	< .001
Controller (-0.5 = Partner, 0.5 = Participant)	6.27	4.02	8.53	6.29	< .001
Training (-0.47 = No Association, 0.53 = Association)	-15.85	-21.29	-10.41	-6.84	< .001
Interaction	13.60	8.91	18.29	6.81	< .001

Note. $N = 432$. *P*-values are based on the Satterthwaite approximation for degrees of freedom. Confidence intervals are based on profile likelihood approximation.

Table 4.2

Growth Model Predicting Mean Number of Tickets Given from Time (Round), Time^2, Controller Assignment, Training Condition, and Their Interactions.

		95% *CI* for *b*			
Predictor	*b*	*LL*	*UL*	*t*	*p*
Intercept	60.16	57.85	62.47	51.17	< .001
Time	5.75	3.69	7.82	5.47	< .001
Time^2	-2.79	-4.06	-1.53	-4.34	< .001
Controller (-0.5 = Partner, 0.5 = Participant)	8.45	6.45	10.45	8.28	< .001
Training (-0.47 = No Association, 0.53 = Association)	-16.99	-21.61	-12.36	-7.21	< .001
*Time*Controller*	25.37	20.49	30.26	10.20	< .001
*Time^2*Controller*	-8.30	-11.31	-5.29	-5.41	< .001
*Time*Training*	-12.75	-16.89	-8.62	-6.06	< .001
*Time^2*Training*	5.30	2.76	7.83	4.11	< .001
*Controller*Training*	19.12	15.11	23.14	9.36	< .001
*Time*Controller*Training*	62.39	52.60	72.18	12.52	< .001
*Time^2*Controller*Training*	-23.22	-29.25	-17.19	-7.57	< .001

Note. $N = 432$. *P*-values are based on the Satterthwaite approximation for degrees of freedom. Confidence intervals are based on profile likelihood approximation. Time and Time^2 are centered orthogonal polynomials representing rounds within Controller assignments.

Table 4.3

Means (Standard Deviations) for Trust Variables, Study 3.

	Low Power	*High Power*	*Overall*
No Association	*n* = 112	*n* = 116	*n* = 228
World Values Survey	2.04 (0.43)	2.00 (0.45)	2.02 (0.44)
Trust Game	604.01 (307.19)	561.80 (309.84)	582.54 (308.58)
Association	*n* = 99	*n* = 105	*n* = 204
World Values Survey	2.08 (0.32)	2.11 (0.48)	2.10 (0.41)
Trust Game	430.64 (302.46)	438.28 (284.96)	434.57 (292.88)
Overall	*n* = 211	*n* = 221	*N* = 432
World Values Survey	2.06 (0.38)	2.05 (0.47)	2.06 (0.42)
Trust Game	522.66 (316.37)	503.11 (303.96)	512.66 (309.88)

Note. World Values Survey is scored on a scale from 1 to 4. Trust Game is scored as the number of tickets (out of 1000) kept by the participant. Higher values indicate greater distrust.

Table 4.4

Linear Models Predicting Trust from Training Condition, Priming Condition, and Their Interaction, Study 3

	World Values Survey					*Trust Game*				
		95% *CI* for *b*					95% *CI* for *b*			
Predictor	*b*	*LL*	*UL*	*t*	*p*	*b*	*LL*	*UL*	*t*	*p*
Training (-0.47 = *No Association,* 0.53 = *Association*)	0.17	-0.01	0.36	1.82	.070	149.03	91.84	206.22	5.12	< .001
Priming (-0.49 = *Low Power,* 0.51 = *High Power*)	-0.03	-0.21	0.16	-0.26	.792	18.67	38.38	75.72	0.64	.520
Interaction	0.17	-0.21	0.55	0.88	.382	-49.85	-164.13	64.44	-0.86	.392
R^2	.01					.06				
F	1.43					9.02				

Note. N = 432. World Values Survey scores standardized before regression.

Table 4.5

Linear Models Predicting Trust from POWER-DISTRUSTING PEOPLE Association, Priming Condition, and Their Interaction, Study 3

	World Values Survey					*Trust Game*				
		95% *CI* for *b*					95% *CI* for *b*			
Predictor	*b*	*LL*	*UL*	*t*	*p*	*b*	*LL*	*UL*	*t*	*p*
POWER-DISTRUSTING PEOPLE Association	0.03	-0.00	0.07	1.72	.086	-1.66	-12.62	9.30	-0.30	.766
Priming (-0.49 = *Low Power,* 0.51 = *High Power*)	-0.04	-0.23	0.15	-0.40	.690	18.29	-40.73	77.30	0.61	.543
Interaction	0.02	-0.05	0.09	0.51	.614	-8.31	-30.26	13.64	-0.74	.457
R^2	.01					.00				
F	1.09					0.33				

Note. N = 432. World Values Survey scores standardized before regression. POWER-DISTRUSTING PEOPLE association centered and scaled so that one unit indicates an increase of 10 points on the 0 (*not at all related*) to 100 (*very related*) scale.

CHAPTER 5: GENERAL DISCUSSION

Social priming research studies the influence of incidentally encountered concepts on later attitudes, judgments, and behavior. This influence is thought to occur through spreading activation of associated concepts (Molden, 2014), leading to the hypothesis that concept associations can be used to predict the presence or absence of social priming effects. This hypothesis was tested in three studies. In Study 1, I obtained estimates of the association strength for 13 concepts pairs from the power priming literature. In Study 2, I used these estimates to predict the presence or absence of power priming effects. And in Study 3, I attempted to strengthen a weak concept association to test the claim that strong concept associations *cause* priming effects.

Study 1 found that concept pairs from the power priming literature vary greatly in the strength of their associations. In a population commonly studied in social priming research, some prime-target pairs (e.g., POWER and BEING IN CHARGE) are strongly associated, while others (e.g., POWER and STEREOTYPING PEOPLE) are weakly associated. If social priming occurs through spreading activation, it should occur for the strongly but not weakly associated concept pairs. In Study 2, participants primed with high versus low power showed significant differences on a strongly associated conceptual dependent variable (TAKING ACTION) but not on moderately (SEEING THINGS FROM OTHERS' PERSPECTIVES) or weakly (DISTRUSTING PEOPLE) associated conceptual dependent variables. These results are consistent with the spreading activation account (e.g., Molden, 2014).

Study 3 attempted to move beyond predicting the presence or absence of priming effects by examining the causal role of concept associations in social priming. I designed an ambitious paradigm intended to manipulate the association between POWER and DISTRUSTING PEOPLE. Participants in Study 3 completed a training task that either paired or did not pair these

concepts. Participants who were exposed to series of trials in which having power was associated with less trustworthy behavior in others exhibited differentially trusting behavior—trusting those who had power over them while distrusting those over whom they had power. These behavioral changes suggest that the manipulation successfully paired these concepts. Still, the training task may have been insufficient to make the dramatic change in concept associations needed to produce a priming effect. Although participants who experienced POWER paired with DISTRUSTING PEOPLE appropriately adjusted their behavior, they did not adjust their explicit association between these concepts and demonstrated no effect of priming POWER on their judgments of others' trustworthiness or on their trusting behavior towards others. These results do not confirm the prediction that manipulating concept associations will affect the ability to observe social priming effects.

Below, I consider limitations of this work and its implications for social priming theory.

5.1 Difficulty of Training Concept Associations

Participants in the Association condition of Study 3 clearly learned the rules of the training game. They showed increasingly distrusting behavior when they had power but increasingly trusting behavior when their partners had power. After 30 rounds, they knew other players behaved in more or less trustworthy ways depending on who had power. Yet they did not adjust their explicit association of POWER with DISTRUSTING PEOPLE. I consider three reasons participants may have adjusted their behavior in the game without reporting a stronger concept association and their implications for the interpretation of Study 3's results.

First, the training may have been insufficiently strong to change associations. Perhaps 30 exposures to an association, lasting about 2.5 minutes total, are not enough. Evidence from the clinical psychology literature suggests that association training may require many more exposures and/or much more time (Hallion & Ruscio, 2011). For example, an intervention to

reduce social anxiety based on 720 trials—more than 20 times the number in Study 3—produced a moderate effect, $d = .47$, on participants' implicit associations of SELF with REJECTION (Clerkin & Teachman, 2010). To put this into context, if participants in the Association condition had reported a mean association a full standard deviation larger than participants in the No Association condition, it would have been only a moderate association of about 60. A strong association would require almost two full standard deviations of change. Such a large shift in concept associations might take hours of training—or perhaps multiple training sessions spread out over several days. If the training were insufficiently strong to change associations, Study 3 leaves open the possibility that spreading activation is responsible for social priming effects but does not speak to them as a the cause of these effects.

Second, the type of training procedure used in Study 3 may affect only task-specific, rather than global, associations. In other words, participants may have learned the rules of the game but have made only very specific inferences from them—that Controller assignment is associated with raffle ticket behavior—and failed to interpret the game in terms of POWER and/or DISTRUSTING PEOPLE at all. If participants did not adjust their POWER-DISTRUSTING PEOPLE association at all, Study 3 leaves open the possibility that concept associations drive social priming effects but does not speak to their causal role.

Third, training may have influenced only implicit, but not explicit, associations. Some models of implicit and explicit attitude change would support this interpretation (e.g., Gawronski & Bodenhausen, 2006)—conditioning will first change implicit associations (perhaps even temporarily) and change explicit associations only after greater exposure. It's possible that the 30 training exposures in Study 3 were enough to shift implicit but not explicit POWER-DISTRUSTING PEOPLE associations. Even though people can report on their implicit associations (Hahn et al., 2014), they may default to reporting their explicit associations unless

specifically asked about implicit ones. Thus, participants in the Association condition of Study 3 may have changed their implicit POWER-DISTRUSTING PEOPLE associations even if they reported the same explicit associations. That no significant interactions between Training condition and Priming condition were found (and that the estimated effects were not consistently in the predicted direction) suggests that any implicit association shift did not affect priming—contrary to the theory of social priming via spreading activation. However, given that changes in implicit associations were not measured, it is impossible to know whether they were sufficiently strong to produce a priming effect or even if they occurred at all.

Whether the training task in Study 3 changed POWER-DISTRUSTING PEOPLE associations at all is uncertain. Thus, although Study 3 failed to find evidence that training condition was related to the strength of concept associations, it leaves open the possibility for a causal role of concept association strength in social priming effects—a possibility that is consistent with the results from the remaining studies in this project. To provide stronger evidence, future research on the role of concept associations in social priming should measure implicit as well as explicit associations and should use more intensive training methods to produce larger shifts in concept associations.

5.2 The Role of Concept Associations in Social Priming

In Studies 1 and 2, and in Studies P1 through P5 in Appendix A, strong but not weak concept associations measured in one sample were associated with statistically significant priming effects in new samples. However, in Study 3, when concept associations and priming effects were estimated on the same participants, priming was not greater for participants with stronger concept associations. What explains this difference, and what does this imply for theories of social priming?

Perhaps the most important consideration in reconciling the findings of Study 3 with the remaining studies is the low power in Study 3 to detect a concept association by Priming condition interaction. Study 3 was adequately powered to detect a large shift in explicit concept associations due to the training task. But no such difference was found. Thus, I reported exploratory models attempting to predict distrust as a function of reported concept associations, Priming condition, and their interaction. However, the skewed distribution of concept associations meant that there were very few participants who reported more than moderate associations between POWER and DISTRUSTING PEOPLE: the 90th percentile on this measure was an association of 71 out of 100. Given that even strong associations show small priming effects (e.g., $d = 0.26$ for POWER and TAKING ACTION in Study 2), observing a statistically significant interaction in this sample is unlikely, even if associations do predict priming effects.

Study 3 represents an ambitious first attempt to manipulate concept associations for the purpose of social priming. Conclusions drawn from it (as with all studies) are limited by power and design features. Yet, the evidence from the remaining studies in this project is consistent with the hypothesis that spreading activation among related constructs is responsible for social priming effects and that their presence and absence can be predicted from these associations. Moreover, models reported in Studies 2 and P5 suggest that social priming occurs only when concept associations exceed a high strength threshold, rather than increasing linearly with association strength.

However, because these studies do not manipulate concept associations, they do not address their causal role in social priming. For this reason, they may also be consistent with other explanations of social priming effects. For example, demand characteristics might also be predictable on the basis of common concept associations: The stronger participants' associations between the conceptual independent and dependent variables, the more they may make

consistent, hypothesis-confirming inferences about their expected behavior. Thus, Studies 1 and 2 (and the studies in Appendix A) are entirely consistent with the spreading activation account of social priming and suggest that concept associations allow one to predict the presence or absence of social priming effects but do not elucidate the causal mechanisms that link associations to priming.

5.3 Implications for Previous Research on Weak Associations

Although I find support for the role of spreading activation among closely related concept in social priming, I also find evidence that many of the concept pairs studied in the power priming literature are moderate or weak in strength. Study 1 found that association strengths for concept pairs studied in the power priming literature vary greatly. These results are paralleled in a preliminary study, Study P1 (see Appendix A), where association strengths for concept pairs from the broader social priming literature were found to cover the entire scale. Moreover, in Study 2, I found support for the hypothesis that only the strongest concept associations are likely to produce social priming effects (see also, Studies P2 through P5 in Appendix A for further support for this claim). Yet, researchers have reported finding priming effects for all of these concept pairs, including the weakly associated ones (e.g., Bargh et al., 1995; Lammers et al., 2009; Gruenfeld et al., 2008). What explains these findings? I consider two possible explanations.

First, priming is not the only cognitive mechanism that gives rise to social judgments and behavior. Perhaps the effects represented by these concept pairs occur for reasons other than spreading activation among related concepts. A typical power priming manipulation, like the one used in Studies 2, 3, P2, and P5, involves vividly recalling and describing a previous experience of high or low power (Galinsky et al., 2003). While this manipulation ought to bring the concept of POWER to mind, it also involves several other cognitive processes, such as recall of memory

and reflection on these experiences. Participants may, while thinking about these experiences, learn from them or reconstrue them, updating their beliefs about what happened. For example, although participants in Study 1 associated POWER with OBJECTIFYING people only weakly, this does not mean that asking Mechanical Turk workers to recall having power wouldn't increase their objectification of others as has been found in previous research (Gruenfeld et al., 2008). Rather, such an effect could occur if recalling powerful situations entailed construing them as situations about meeting goals (consider the strong association between POWER and PURSUING GOALS in Study 1), thus recasting low power people as instruments that can help or hinder goal pursuit. This recasting could occur even if people do not hold a concept association between POWER and OBJECTIFYING OTHERS. Thus, it is important to rule out alternative mechanisms or to measure concept associations to establish the role of spreading activation in priming effects.

Second, differences in concept associations between populations or samples may lead to differences in priming effects. For example, although the association for POWER and SEX in Study 1 was weak for the whole sample considered together (*median* = 23.5), it differed substantially between men (*median* = 45) and women (*median* = 9). Moreover, men varied a great deal on this association (*median absolute deviation* = 51.89), whereas women were comparatively homogeneous (*median absolute deviation* = 13.34). Early work on POWER-SEX priming used only male participants (Bargh et al., 1995). It would be entirely consistent with the data reported here if the men in these studies had strong associations between POWER and SEX and thus the priming results reported hold true. The associations reported here hold for Mechanical Turk workers, but they may differ in other populations. Thus, they cannot speak to the role of spreading activation in priming effects obtained in other populations.

5.4 Effect Sizes in Social Priming

Despite the uncertainty in the interpretation of Study 3, there is evidence in Study 2 and in Studies P2 through P5 in Appendix A that social priming effects occur for strong concept association. The studies reported here are some of the largest social priming studies ever conducted and thus provide some of the best evidence about the presence, absence, and size of these effects. Examining the effect size estimates from these studies leads to the conclusion that social priming effects are much smaller than in many initial reports, and may be only weakly detectable in typical sample sizes in this field.

A recent meta-analysis of prime-to-behavior effects found an average effect size of d = 0.33 and an average sample size of fewer than 25 participants per condition (Weingarten et al., 2016). With one exception, all of the effect sizes from the studies reported here are smaller than this estimate, and all of the sample sizes far exceed 25 participants per condition—usually by a factor of 10 or more. Figure 5.1 (page 48) presents a forest plot of 19 effect sizes from Studies P2, P3, P4, P5, 2, and 3 (Studies P2 through P5 are reported in Appendix A). Most of the effects investigated as part of the project are approximately 0. Their overall median is d = 0.03. Those effects that differ from 0, using an uncorrected α of .05, occur almost exclusively for large associations and range in absolute value from d = 0.17 to d = 0.43. Their median is d = 0.25. A study that had 80% power to detect an effect of this size in a two-group between subjects design would have 252 participants per condition. 25 subjects per condition would provide only 14% power to detect an effect of this size.

Thus, future studies of social priming effects should recruit much larger samples than earlier studies of these phenomena—at least 250 participants per condition. With such large sample sizes, the value of measuring concept associations becomes starkly clear. Concept associations that usefully predict the presence or absence of social priming effects are easy and

fast to collect using explicit measures like those in Studies 1 and P1. Before recruiting 500 or many hundreds of participants for a social priming study, researchers can make reasonable judgments about the value of such an investment by asking people how strongly they associate the concepts under study.

5.5 Conclusion

The correlation between association strength and effect size in the 19 priming effects from this project is $r(17) = .41$ [-.06, .73], $p = .08$.[5] Although not significant, the relationship is positive and consistent with larger associations predicting larger priming effects—or with a threshold model in which only the largest associations produce social priming effects. And, in general, significant effects are found only for the strongest concept associations. Importantly, these results do not imply that reported social priming effects based on weak concept associations are false positives. Rather, such effects may have been discovered in populations where these concept associations are actually strong or may rely instead on mechanisms other than concept accessibility, such as learning or persuasion. Future research could more firmly clarify the causal role of concept accessibility in social priming by using more intensive manipulations of concept associations, measuring implicit as well as explicit associations, and using sample sizes of more than 250 participants per condition.

[5] A rank correlation produces a smaller estimate of this relationship, $\tau_B(17) = .22$, $p = .21$.

5.6 Figures

study	association	strength	effect	N
P2	power - being in charge	96	Sentence Unscrambling	779
P2	power - being in charge	96	Recall	691
P3	time - self-reflection	21	Sentence Unscrambling	767
P3	time - self-reflection	21	Recall	683
P4	god - watching	41	Public Self-Awareness	634
P4	god - watching	41	Private Self-Awareness	634
P4	god - watching	41	Negative Self-Conscious Emotions	634
P4	god - watching	41	Pride	634
P5	power - being in charge	96	Being in Charge	521
P5	power - self-esteem	47	Self-Esteem	521
P5	power - distrust	21	WVS Trust	521
P5	power - distrust	21	Behavioral Trust	521
2	power - taking action	87	Action/Inaction Goals	579
2	power - seeing things from others' perspectives	51	IRI Perspective-Taking	579
2	power - seeing things from others' perspectives	51	Reading the Mind in the Eyes Test	579
2	power - distrusting people	28	WVS Trust	579
2	power - distrusting people	28	Trust Game	579
3	power - distrusting people	28	WVS Trust	432
3	power - distrusting people	28	Trust Game	432

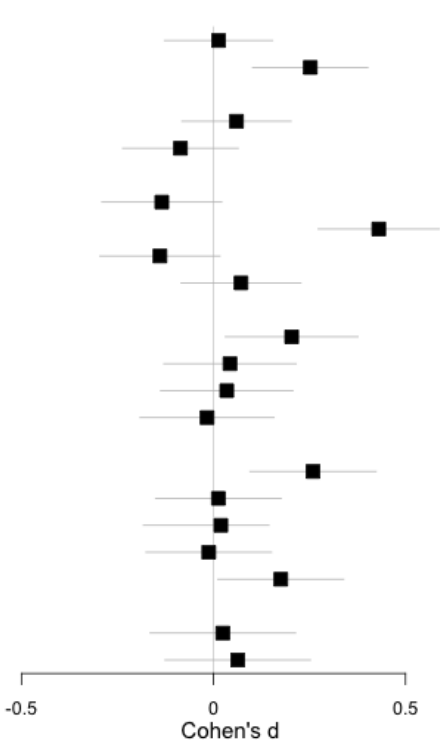

Figure 5.1. Forest plot showing effect sizes and 95% confidence intervals for 20 social priming effects. Studies P2 through P5 are reported in Appendix A. All effects are keyed so that hypothesized results are positive. In Studies P2 and P3, participants were randomly assigned to one of two priming methods (Sentence Unscrambling or Recall), analyzed separately here. In Studies P3, P4, P5, 2, and 3, some conceptual dependent variables were measured with multiple weakly correlated indicators, each of which is analyzed separately here. Concept association strengths were measured in Study P1 (Appendix A) for effects in Studies P2, P3, P4, and P5; in Study 1 for effects in Study 2; and in Study 3 for effects in Study 3. WVS = World Values Survey. IRI = Interpersonal Reactivity Index.

REFERENCES

Adler, N. E., Boyce, T., Chesney, M. A., Cohen, S., Folkman, S., Kahn, R. L., & Syme, S. L. (1994). Socioeconomic status and health: the challenge of the gradient. *American Psychologist, 24*. http://dx.doi.org/10.1037/0003-066X.49.1.15

Albarracín, D. (2009). *Action/inaction goal scales*. Unpublished measures. Department of Psychology, University of Illinois, Champaign, IL.

Anderson, C., Srivastava, S., Beer, J., Spataro, S. E., & Chatman, J. A. (2006). Knowing your place: Self-perceptions of status in social groups. *Journal of Personality and Social Psychology, 91*, 1094–1110. http://dx.doi.org/10.1037/0022-3514.91.6.1094

Arndt, J., Greenberg, J., Pyszczynski, T., & Solomon, S. (1997). Subliminal exposure to death-related stimuli increases defense of the cultural worldview. *Psychological Science, 8*, 379–385. http://dx.doi.org/10.1111/j.1467-9280.1997.tb00429.x

Baldwin, M. W., & Holmes, J. G. (1987). Salient private audiences and awareness of the self. *Journal of Personality and Social Psychology, 52*, 1087-1098. http://dx.doi.org/10.1037/0022-3514.52.6.1087

Bargh, J. A., Bond, R. N., Lombardi, W. J., & Tota, M. E. (1986). The additive nature of chronic and temporary sources of construct accessibility. *Journal of Personality and Social Pscyhology, 50*, 869–878. http://dx.doi.org/10.1037/0022-3514.50.5.869

Bargh, J. A., Chaiken, S., Govender, R., & Pratto, F. (1992). The generality of the automatic attitude activation effect. *Journal of Personality and Social Psychology, 62*, 893–912. http://dx.doi.org/10.1037/0022-3514.62.6.893

Bargh, J. A., Chen, M., & Burrows, L. (1996). Automaticity of social behavior: Direct effects of trait construct and stereotype activation on action. *Journal of Personality and Social Psychology, 71*, 230–244. http://dx.doi.org/10.1037/0022-3514.71.2.230

Bargh, J. A., Raymond, P., Pryor, J. B., & Strack, F. (1995). Attractiveness of the underling: an automatic power → sex association and its consequences for sexual harassment and aggression. *Journal of Personality and Social Psychology, 68*, 768–781. http://dx.doi.org/10.1037/0022-3514.68.5.768

Baron-Cohen, S., Wheelwright, S., Hill, J., Raste, Y., & Plumb, I. (2001). The "Reading the Mind in the Eyes" Test revised version: A study with normal adults, and adults with Asperger syndrome or high-functioning autism. *Journal of Child Psychology and Psychiatry, 42*, 241–251. http://dx.doi.org/10.1111/1469-7610.00715

Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior, 10*, 122–142. http://dx.doi.org/10.1006/game.1995.1027

Bering, J. M., & Johnson, D. D. P. (2005). "O Lord… you perceive my thoughts from afar": Recursiveness and the evolution of supernatural agency. *Journal of Cognition and Culture, 5*, 118–142. http://dx.doi.org/10.1163/1568537054068679

Blackie, L. E. R., & Cozzolino, P. J. (2011). Of blood and death: A test of dual-existential systems in the context of prosocial intentions. *Psychological Science, 22*, 998–1000. http://dx.doi.org/10.1177/0956797611415542

Brandt, M. J., IJzerman, H., & Blanken, I. (2014). Does recalling moral behavior change the perception of brightness? A replication and meta-analysis of Banerjee, Chatterjee, and Sinha (2012). *Social Psychology, 45*, 246–252. http://dx.doi.org/10.1027/1864-9335/a000191

Briñol, P., Petty, R. E., Valle, C., Rucker, D. D., & Becerra, A. (2007). The effects of message recipients' power before and after persuasion: A self-validation analysis. *Journal of Personality and Social Psychology, 93*, 1040–1053. http://dx.doi.org/10.1037/0022-3514.93.6.1040

Carter, T. J., Ferguson, M. J., & Hassin, R. R. (2011). A single exposure to the American flag shifts support toward Republicanism up to 8 months later. *Psychological Science, 22*, 1011–1018. http://dx.doi.org/10.1177/0956797611414726

Cattell, J. M. (1887). Experiments on the association of ideas. *Mind, 12*, 68–74. http://dx.doi.org/10.1093/mind/os-12.45.68

Cesario, J., & Jonas, K. J. (2014). Replicability and models of priming: What a resource computation framework can tell us about expectations of replicability. *Social Cognition*, *32*(Supplement), 124–136. http://doi.org/10.1521/soco.2014.32.supp.124

Cesario, J., Plaks, J. E., Hagiwara, N., Navarrete, C. D., & Higgins, E. T. (2010). The ecology of automaticity: How situational contingencies shape action semantics and social behavior. *Psychological Science, 21*, 1311–1317. http://dx.doi.org/10.1177/0956797610378685

Cesario, J., Plaks, J. E., & Higgins, E. T. (2006). Automatic social behavior as motivated preparation to interact. *Journal of Personality and Social Psychology, 90*, 893–910. http://dx.doi.org/10.1037/0022-3514.90.6.893

Chen, S., Lee-Chai, A. Y., & Bargh, J. A. (2001). Relationship orientation as a moderator of the effects of social power. *Journal of Personality and Social Psychology, 80*, 173–187. http://dx.doi.org/10.1037/0022-3514.80.2.173

Clerkin, E. M., & Teachman, B. A. (2010). Training implicit social anxiety associations: An experimental intervention. *Journal of Anxiety Disorders, 24*, 300–308. http://dx.doi.org/10.1016/j.janxdis.2010.01.001

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review, 82*, 407–428. http://dx.doi.org/10.1037/0033-295X.82.6.407

Davis, M. H. (1980). A multidimensional approach to individual differences in empathy. JSAS *Catalog of Selected Documents in Psychology, 10*, 85.

Doyen, S., Klein, O., Pichon, C.-L., & Cleeremans, A. (2012). Behavioral priming: It's all in the mind, but whose mind? *PLoS ONE, 7*, e29081. http://dx.doi.org/10.1371/journal.pone.0029081

Doyen, S., Klein, O., Simons, D., & Cleeremans, A. (2014). On the other side of the mirror. Priming in cognitive and social psychology. *Social Cognition, 32(Supplement)*, 12–32. http://dx.doi.org/10.1521/soco.2014.32.supp.12

Dubois, D. (2011). *The power matching effect: Why the powerful persuade the powerful but the powerless persuade the powerless* (Doctoral dissertation). Retrieved from ProQuest Dissertations & Theses. (3469713)

Earp, B. D., Everett, J. A. C., Madva, E. N., & Hamlin, J. K. (2014). Out, damned spot: Can the "Macbeth effect" be replicated? *Basic and Applied Social Psychology, 36*, 91–98. http://dx.doi.org/10.1080/01973533.2013.856792

Eskine, K. J., Kacinik, N. A., & Prinz, J. J. (2011). A bad taste in the mouth: Gustatory disgust influences moral judgment. *Psychological Science, 22*, 295–299. http://dx.doi.org/10.1177/0956797611398497

Fast, N. J., Gruenfeld, D. H., Sivanathan, N., & Galinsky, A. D. (2009). Illusory control: A generative force behind power's far-reaching effects. *Psychological Science, 20*, 502–508. http://dx.doi.org/10.1111/j.1467-9280.2009.02311.x

Fayard, J. V., Bassi, A. K., Bernstein, D. M., & Roberts, B. W. (2009). Is cleanliness next to godliness? Dispelling old wives' tales: Failure to replicate Zhong and Liljenquist (2006). *Journal of Articles in Support of the Null Hypothesis, 6*, 21–30.

Fazio, R. H. (1993). Variability in the likelihood of automatic attitude activation: Data reanalysis and commentary on Bargh, Chaiken, Govender, and Pratto (1992). *Journal of Personality and Social Psychology, 64*, 753–758. http://dx.doi.org/10.1037/0022-3514.64.5.753

Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology, 50*, 229–238. http://dx.doi.org/10.1037/0022-3514.50.2.229

Ferguson, M. J., Carter, T. J., & Hassin, R. R. (2014). Commentary on the attempt to replicate the effect of the American flag on increased republican attitudes. *Social Psychology, 45*, 301-302. http://dx.doi.org/10.1027/1864-9335/a000202

Fischler, I. (1977). Associative facilitation without expectancy in a lexical decision task. *Journal of Experimental Psychology: Human Perception and Performance*, *3*, 18–26. http://dx.doi.org/10.1037/0096-1523.3.1.18

Galinsky, A. D., Gruenfeld, D. H., & Magee, J. C. (2003). From power to action. *Journal of Personality and Social Psychology, 85*, 453–466. http://dx.doi.org/10.1037/0022-3514.85.3.453

Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin, 132*, 692–731. http://dx.doi.org/10.1037/0033-2909.132.5.692

Gervais, W. M., & Norenzayan, A. (2012). Like a camera in the sky? Thinking about God increases public self-awareness and socially desirable responding. *Journal of Experimental Social Psychology, 48*, 298–302. http://dx.doi.org/10.1016/j.jesp.2011.09.006

Gillath, O., Mikulincer, M., Birnbaum, G. E., & Shaver, P. R. (2008). When sex primes love: Subliminal sexual priming motivates relationship goal pursuit. *Personality and Social Psychology Bulletin, 34*, 1057–1069. http://dx.doi.org/10.1177/0146167208318141

Gino, F., & Mogilner, C. (2014). Time, money, and morality. *Psychological Science, 25*, 414–421. http://dx.doi.org/10.1177/0956797613506438

Govern, J. M., & Marsch, L. A. (2001). Development and validation of the situational self-awareness scale. *Consciousness and Cognition, 10*, 366–378. http://dx.doi.org/10.1006/ccog.2001.0506

Grant, A. M., Franklin, J., & Langford, P. (2002). The self-reflection and insight scale: A new measure of private self-consciousness. *Social Behavior and Personality: an International Journal, 30*, 821–835. http://dx.doi.org/10.2224/sbp.2002.30.8.821

Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review, 102*, 4–27. http://dx.doi.org/10.1037/0033-295X.102.1.4

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: the implicit association test. *Journal of Personality and Social Psychology, 74*, 1464–1480. http://dx.doi.org/10.1037/0022-3514.74.6.1464

Gruenfeld, D. H., Inesi, M. E., Magee, J. C., & Galinsky, A. D. (2008). Power and the objectification of social targets. *Journal of Personality and Social Psychology, 95*, 111–127. http://dx.doi.org/10.1037/0022-3514.95.1.111

Guinote, A. (2007). Power and goal pursuit. *Personality and Social Psychology Bulletin, 33*, 1076–1087. http://dx.doi.org/10.1177/0146167207301011

Hahn, A., Judd, C. M., Hirsh, H. K., & Blair, I. V. (2014). Awareness of implicit attitudes. *Journal of Experimental Psychology: General, 143*, 1369–1392. http://dx.doi.org/10.1037/a0035028

Hallion, L. S., & Ruscio, A. M. (2011). A meta-analysis of the effect of cognitive bias modification on anxiety and depression. *Psychological Bulletin, 137*, 940–958. http://dx.doi.org/10.1037/a0024355

Harris, C. R., Coburn, N., Rohrer, D., & Pashler, H. (2013). Two failures to replicate high-performance-goal priming effects. *PLoS ONE, 8*, e72467. http://dx.doi.org/10.1371/journal.pone.0072467

Heatherton, T. F., & Polivy, J. (1991). Development and validation of a scale for measuring state self-esteem. *Journal of Personality and Social Psychology, 60*, 895–910. http://dx.doi.org/10.1037/0022-3514.60.6.895

Higgins, E. T., Rholes, W. S., & Jones, C. R. (1977). Category accessibility and impression formation. *Journal of Experimental Social Psychology, 13*, 141–154. http://dx.doi.org/10.1016/S0022-1031(77)80007-3

Hirschberger, G., Ein-Dor, T., & Almakias, S. (2008). The self-protective altruist: Terror management and the ambivalent nature of prosocial behavior. *Personality and Social Psychology Bulletin, 34*, 666–678. http://dx.doi.org/10.1177/0146167207313933

Inagaki, T. K., & Eisenberger, N. I. (2013). Shared neural mechanisms underlying social warmth and physical warmth. *Psychological Science, 24*, 2272–2280. http://dx.doi.org/10.1177/0956797613492773

Inesi, M. E., Gruenfeld, D. H., & Galinsky, A. D. (2012). How power corrupts relationships: Cynical attributions for others' generous acts. *Journal of Experimental Social Psychology, 48*, 795–803. http://dx.doi.org/10.1016/j.jesp.2012.01.008

Johnson, D., & Bering, J. M. (2006). Hand of God, mind of man: punishment and cognition in the evolution of cooperation. *Evolutionary Psychology, 4*, 219–233. http://dx.doi.org/10.1177/147470490600400119

Johnson, D. J., Cheung, F., & Donnellan, M. B. (2014). Does cleanliness influence moral judgments? A direct replication of Schnall, Benton, and Harvey (2008). *Social Psychology, 45*, 209–215. http://dx.doi.org/10.1027/1864-9335/a000186

Jonas, E., Schimel, J., Greenberg, J., & Pyszczynski, T. (2002). The Scrooge effect: Evidence that mortality salience increases prosocial attitudes and behavior. *Personality and Social Psychology Bulletin, 28*, 1342–1353. http://dx.doi.org/10.1177/014616702236834

Kahneman, D. (2012, September 6). A proposal to deal with questions about social priming effects [open letter]. *Nature*. Retrieved from http://www.nature.com/polopoly_fs/7.6716.1349271308!/suppinfoFile/Kahneman%20Letter.pdf

Kasser, T., & Sheldon, K. M. (2000). Of wealth and death: Materialism, mortality salience, and consumption behavior. *Psychological Science, 11*, 348–351. http://dx.doi.org/10.1111/1467-9280.00269

Kernis, M. H., & Goldman, B. M. (2006). A multicomponent conceptualization of authenticity: Theory and research. *Advances in Experimental Social Psychology, 38*, 283–357. http://dx.doi.org/10.1016/S0065-2601(06)38006-9

Kitchens, M. B. (2015). Thinking about God causes internal reflection in believers and unbelievers of God. *Self and Identity, 14*, 724-747. http://dx.doi.org/10.1080/15298868.2015.1072108

Klein, R. A., Ratliff, K. A., Vianello, M., Adams, R. B., Jr., Bahník, Š., Bernstein, M. J., … Nosek, B. A. (2014). Investigating variation in replicability: A "many labs" replication project. *Social Psychology, 45*, 142–152. http://dx.doi.org/10.1027/1864-9335/a000178

Kraus, M. W., Piff, P. K., & Keltner, D. (2011). Social class as culture: The convergence of resources and rank in the social realm. *Current Directions in Psychological Science, 20*, 246–250. doi:10.1177/0963721411414654

Lammers, J., Dubois, D., Rucker, D. D., & Galinsky, A. D. (2013). Power gets the job: Priming power improves interview outcomes. *Journal of Experimental Social Psychology, 49*, 776–779. http://dx.doi.org/10.1016/j.jesp.2013.02.008

Lammers, J., Stoker, J. I., & Stapel, D. A. (2009). Differentiating social and personal power: Opposite effects on stereotyping, but parallel effects on behavioral approach tendencies. *Psychological Science, 20*, 1543–1549. http://dx.doi.org/10.1111/j.1467-9280.2009.02479.x

Loersch, C., & Payne, B. K. (2011). The situated inference model: An integrative account of the effects of primes on perception, behavior, and motivation. *Perspectives on Psychological Science, 6*, 234–252. http://dx.doi.org/10.1177/1745691611406921

Loersch, C., & Payne, B. K. (2014). Situated inferences and the what, who, and where of priming. *Social Cognition, 32*, 137–151. http://dx.doi.org/10.1521/soco.2014.32.supp.137

McCullough, M. E., Carter, E. C., DeWall, C. N., & Corrales, C. M. (2012). Religious cognition down-regulates sexually selected, characteristically male behaviors in men, but not in women. *Evolution and Human Behavior, 33*, 562–568. http://dx.doi.org/10.1016/j.evolhumbehav.2012.02.004

Meuller, P., Chandler, J., & Paolacci, G. (2012). *Advanced uses of Mechanical Turk in psychological research*. Paper presented at the Society for Personality and Social Psychology, San Diego, CA. Slides retrieved from https://experimentalturk.files.wordpress.com/2012/01/mueller-spsp-2012.pdf

Meyer, D. 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. http://dx.doi.org/10.1037/h0031564

Molden, D. C. (2014). Understanding priming effects in social psychology: What is "social priming" and how does it occur? *Social Cognition, 32*(Supplement), 1–11. http://doi.org/10.1521/soco.2014.32.supp.1

Mortensen, C. R., Ackerman, J. M., Becker, D. V., Neuberg, S. L., & Kenrick, D. T. (2010). Infection breeds reticence: The effects of disease salience on self-perceptions of personality and behavioral avoidance tendencies. *Psychological Science, 21*, 440–447. http://dx.doi.org/10.1177/0956797610361706

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. http://dx.doi.org/10.1037/0096-3445.106.3.226

Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (2004). The University of South Florida free association, rhyme, and word fragment norms. *Behavior Research Methods, Instruments, & Computers, 36*, 402–407. http://dx.doi.org/10.3758/BF03195588

Pashler, H., Coburn, N., & Harris, C. R. (2012). Priming of social distance? Failure to replicate effects on social and food judgments. *PLoS ONE, 7*, e42510. http://dx.doi.org/10.1371/journal.pone.0042510

Pavlov, I. P. (2003). *Conditioned reflexes*. Mineola, NY: Dover Publications. (Original work published 1927)

Porter, L. W. (1961). A study of perceived need satisfactions in bottom and middle management jobs. *Journal of Applied Psychology, 45*, 1–10. http://dx.doi.org/10.1037/h0043121

Robins, R. W., & Hendin, H. M. (2001). Measuring global self-esteem: Construct validation of a single-item measure and the Rosenberg Self-Esteem Scale. *Personality and Social Psychology Bulletin, 27*, 151–161. http://dx.doi.org/10.1177/0146167201272002

Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology, 7*, 573–605. http://dx.doi.org/10.1016/0010-0285(75)90024-9

Rosenberg, M. (1965). *Society and the adolescent self-image*. Princeton, NJ: Princeton University Press.

Royston, P., Altman, D. G., & Sauerbrei, W. (2005). Dichotomizing continuous predictors in multiple regression: A bad idea. *Statistics in Medicine*, *25*, 127–141. http://dx.doi.org/10.1002/sim.2331

Ryff, C. D., & Keyes, C. L. M. (1995). The structure of psychological well-being revisited. *Journal of Personality and Social Psychology, 69*, 719–727. http://dx.doi.org/10.1037/0022-3514.69.4.719

Schmid Mast, M., Jonas, K., & Hall, J. A. (2009). Give a person power and he or she will show interpersonal sensitivity: The phenomenon and its why and when. *Journal of Personality and Social Psychology, 97*, 835–850. http://dx.doi.org/10.1037/a0016234

Segal, S. J. (1967). Priming of association test responses by differential verbal contexts. *Journal of Experimental Psychology, 74*, 370–377. http://dx.doi.org/10.1037/h0024704

Shariff, A., Cohen, A., & Norenzayan, A. (2008). The devil's advocate: Secular arguments diminish both implicit and explicit religious belief. *Journal of Cognition and Culture, 8*, 417–423. http://dx.doi.org/10.1163/156853708X358245

Shariff, A., & Norenzayan, A. (2007). God is watching you: Priming God concepts increases prosocial behavior in an anonymous economic game. *Psychological Science, 18*, 803–809. http://10.1111/j.1467-9280.2007.01983.x

Shariff, A., Willard, A. K., Andersen, T., & Norenzayan, A. (2016). Religious priming: A meta-analysis with a focus on prosociality. *Personality and Social Psychology Review, 20*, 27-28. http://dx.doi.org/10.1177/1088868314568811

Sherman, G. D., Haidt, J., & Clore, G. L. (2012). The faintest speck of dirt: Disgust enhances the detection of impurity. *Psychological Science, 23*, 1506–1514. http://dx.doi.org/10.1177/0956797612445318

Smith, P. K., & Bargh, J. A. (2008). Nonconscious effects of power on basic approach and avoidance tendencies. *Social Cognition, 26*, 1-24. http://dx.doi.org/10.1521/soco.2008.26.1.1

Smith, P. K., Dijksterhuis, A., & Wigboldus, D. H. J. (2008). Powerful people make good decisions even when they consciously think. *Psychological Science, 19*, 1258–1259. http://doi.org/10.1111/j.1467-9280.2008.02207.x

Smith, P. K., Jostmann, N. B., Galinsky, A. D., & van Dijk, W. W. (2008). Lacking power impairs executive functions. *Psychological Science, 19*, 441–447. http://doi.org/10.1111/j.1467-9280.2008.02107.x

Smith, P. K., & Trope, Y. (2006). You focus on the forest when you're in charge of the trees: Power priming and abstract information processing. *Journal of Personality and Social Psychology, 90*, 578–596. http://dx.doi.org/10.1037/0022-3514.90.4.578

Srull, T. K., & Wyer, R. S., Jr. (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. http://dx.doi.org/10.1037/0022-3514.37.10.1660

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology, 5*, 207–232. http://dx.doi.org/10.1016/0010-0285(73)90033-9

Van Lange, P. A., Otten, W., De Bruin, E. M., & Joireman, J. A. (1997). Development of prosocial, individualistic, and competitive orientations: theory and preliminary evidence. *Journal of Personality and Social Psychology, 73*, 733–746. http://dx.doi.org/10.1037/0022-3514.73.4.733

Weingarten, E., Chen, Q., McAdams, M., Yi, J., Hepler, J., & Albarracín, D. (2016). From primed concepts to action: A meta-analysis of the behavioral effects of incidentally presented words. *Psychological Bulletin, 142*, 472-497. http://doi.org/10.1037/bul0000030

Wentura, D., & Rothermund, K. (2014). Priming is not priming is not priming. *Social Cognition, 32,* 47–67. http://dx.doi.org/10.1521/soco.2014.32.supp.47

Wheeler, S. C., DeMarree, K. G., & Petty, R. E. (2007). Understanding the role of the self in prime-to-behavior effects: The active-self account. *Personality and Social Psychology Review, 11*, 234–261. http://dx.doi.org/10.1177/1088868307302223

Wheeler, S. C., DeMarree, K. G., & Petty, R. E. (2014). Understanding prime-to-behavior effects: Insights from the active-self account. *Social Cognition, 32*, 109–123. http://dx.doi.org/10.1521/soco.2014.32.supp.109

Williams, L. E. (2014). Improving psychological science requires theory, data, and caution: Reflections on Lynott et al. (2014). *Social Psychology, 45,* 321-326. http://dx.doi.org/10.1027/1864-9335/a000205

Williams, L. E., & Bargh, J. A. (2008). Experiencing physical warmth promotes interpersonal warmth. *Science, 322*, 606–607. http://dx.doi.org/10.1126/science.1162548

World Values Survey Association. (2012). *2010 -2012 World Values Survey*. Madrid, Spain: Asep/JDS. Retrieved from http://www.worldvaluessurvey.org/WVSDocumentationWV6.jsp

APPENDIX A: PRELIMINARY STUDIES P1 THROUGH P5

Five preliminary studies were run as part of this project. Data for these studies were collected before the three studies reported in the main text. Study P1 aimed to collect estimates of population-level association strengths for concept pairs from the priming literature. It differed from Study 1 in that it sampled associations for ten different prime concepts (e.g., associations with GOD and DEATH in addition to POWER). Each of these prime concepts was presented with one or two target concepts from the priming literature (e.g., HOLDING SOMETHING HOT with INTERPERSONAL WARMTH) as well as several distractor concepts (e.g., HOLDING SOMETHING HOT with APPLE). From the 13 prime-target pairs in Study P1, I selected a strong and weak concept association pair to test the hypothesis that priming would occur for the strong association but not for the weak one. Because these pairs were based on different prime concepts, their effects were tested in separate studies—Studies P2 (strong association) and P3 (weak association). Study P4 tested the related prediction that social priming effects for a given concept pair (GOD and WATCHING) will differ for populations with strong (religious Mechanical Turk workers) vs. weak (non-religious Mechanical Turk workers) associations. Finally, Study P5 used a method similar to Study 2, measuring the effect of a single prime concept on three target conceptual dependent variables that varied in association strength (two target associations and a distractor association from Study P1). If concept association strength is predictive of priming effects, priming should be observed for the strong but not weak concept associations, both across concept pairs (i.e., between strong and weak pairs in Studies P2, P3, and P5) and within concept pairs across individuals who differ in ways that affect the strength of the association (i.e., between individuals in Study P4).

A.1 Study P1

Study P1 had two goals: (1) to estimate the strength of association for concept pairs from

the social priming literature in a population commonly used for priming research, workers on Amazon's Mechanical Turk service, and (2) to probe for individual differences that predict differences in concept association strengths. To estimate association strength, participants were asked to rate how strongly thirteen concept pairs from the priming literature were related. These association strengths would then be used to inform selection of associations for later studies.

A.1.1 Method

Participants. To gain accurate estimates of association strength, I aimed to collect data from 700 participants after exclusions for non-US IP addresses and failing an instructional manipulation check. To meet this goal, I recruited 806 mechanical Turk workers, who were paid $0.80 each for participation. Eighteen workers were excluded for non-US IP addresses, and 38 were excluded for failing an instructional manipulation check (including 2 excluded for both reasons), leaving 752 workers (268 women; 484 men) included in analyses. They averaged 31.84 years old (*SD* = 10.89), and most identified as European American (77%). Most reported being non-religious (56%), with Christians making up the next largest group (38%). Most participants (57%) had completed at least some college, and 57% reported an annual income of $35,000 or less.

Procedure. Participants completed ten blocks of self-report concept association questions like that in Study 1. Each block featured a different focal word or phrase (a prime from the literature). Participants judged how closely the focal concept was related to five or six different concepts, using a scale ranging from 0 (*not at all related*) to 100 (*strongly related*). For each focal word, one or two of the judgments concerned target concepts from the priming literature, and the remaining four judgments concerned distractor concepts. Distractor concepts were included to encourage participants to use the whole scale, to obscure the target concepts, and to validate our measure of concept associations. For each focal concept, the distractors represented

two concepts (one concrete and one abstract) expected to have a strong association with the focal concept and two concepts (one concrete and one abstract) expected to have a weak association with the focal concept. The full list of distractor concepts is presented in Table A.1 (page 84). The order of blocks and the order of items within blocks were randomized. The thirteen prime-target concept pairs assessed were:

- POWER and BEING IN CHARGE (Fast et al., 2009)
- POWER and SELF-ESTEEM (Briñol et al., 2007)
- ELDERLY PEOPLE and SLOW WALKING (Bargh et al., 1996; Cesario, Plaks, & Higgins, 2006)
- CLEANLINESS and WHITE (Sherman, Haidt, & Clore, 2012)
- PATRIOTISM and CONSERVATISM (Carter et al., 2011)
- GOD and WATCHING (Gervais & Norenzayan, 2012)
- GOD and GENEROSITY (Shariff & Norenzayan, 2007)
- SEXY PICTURES and RELATIONSHIP CLOSENESS (Gillath, Mikulincer, Birnbaum, & Shaver, 2008)
- TIME and SELF-REFLECTION (Gino & Mogilner, 2014)
- SUGARY and KINDNESS (Eskine, Kacinik, & Prinz, 2011)
- HOLDING SOMETHING HOT and INTERPERSONAL WARMTH (Inagaki & Eisenberger, 2013; Williams & Bargh, 2008)
- DEATH and SELFISHNESS (Kasser & Sheldon, 2000)
- DEATH and GENEROSITY (Blackie & Cozzolino, 2011; Hirschberger, Ein-Dor, & Almakias, 2008; Jonas, Schimel, Greenberg, & Pyszczynski, 2002).

Participants also completed the same demographics questionnaire used in Study 1.

A.1.2 Results

Because several of the concept association distributions were highly skewed (skewness range: $-4.22 \leq g_1 \leq 6.28$), medians were used as measures of central tendency. Using means does not substantially change the interpretation of results, though it does restrict the range of central tendencies (see Appendix B for a figure comparing median and mean association strengths for the target items).

Distractor items. The distractor items allow me to evaluate the validity of this measure of concept associations. These ratings for these items generally conformed to expectations. The median association strength for *concrete related* distractors ranged from 36 (PATRIOTISM and RED) to 99 (TIME and CLOCK). Although the association strength for PATRIOTISM and RED was lower than expected, this pair appears to be unusual among the concrete unrelated distractors, with the remaining median ratings all being above the midpoint of 50. Median ratings for *abstract related* distractors were very high, ranging from 81 (POWER and STATUS) to 100 (ELDERLY PEOPLE and OLD; GOD and LORD; DEATH and DECEASED). As expected, the two types of unrelated distractors received comparatively low association strength ratings, with all medians for *concrete unrelated* distractors and for seven of the *abstract unrelated* distractors equal to 0. The remaining *abstract unrelated* distractors had median association strength ratings of 1 (CLEANLINESS and INTELLIGENCE), 9 (ELDERLY PEOPLE and ORGANIZED), and 22 (POWER and DISTRUST). Overall, the judgments for distractor items suggest that this self-report method is a valid measure of concept associations.

Target items. Figure A.1a (page 90) depicts the median concept association scores for the 13 prime-target pairs from the social psychological literature included in this study. Mechanical Turk workers' reported strength of associations varied greatly, ranging from 0 (DEATH and GENEROSITY; DEATH and SELFISHNESS) to 96 (POWER and BEING IN

CHARGE). Figure A.1a also shows substantial variability in the heterogeneity of participants' responses to the 13 target items, indicated by differences in the shapes of their distributions and the widths of the confidence intervals. Mechanical Turk workers are relatively more similar to each other in how strongly they associate some concept pairs, such as DEATH and GENEROSITY (*CI* undefined), and relatively less similar to each other in how strongly they associate other concept pairs, such as GOD and GENEROSITY (*CI width* = 10).

Individual differences. Six demographic variables—gender, religious affiliation, age, political conservatism, objective SES, and subjective SES—were used to explore whether differences in concept associations can be predicted from individual differences expected to influence social cognition.

Gender. Bias-corrected and accelerated confidence intervals for the difference in medians between men and women were estimated using 10,000 bootstrapped resamplings of the data. Only one of these confidence intervals excluded zero: women reported a stronger association between ELDERLY PEOPLE and SLOW WALKING than did men (*difference in medians* = 6 [1, 9]) (see Figure A.1b).

Religion. Participants' responses to the question about religious affiliation were coded into two categories, religious (e.g., "Baptist") or non-religious (e.g. "agnostic"). Bias corrected and accelerated confidence intervals around the differences in medians between these groups were again estimated using 10,000 bootstrapped resamplings. Three of these confidence intervals excluded zero. Religious participants reported stronger associations than did non-religious participants for CLEANLINESS and WHITE (*difference in medians* = 9 [1, 11]), GOD and WATCHING (*difference in medians* = 11 [1, 18]), and GOD and GENEROSITY (*difference in medians* = 3.5 [23, 39]) (see Figure A.1c).

Other demographic variables. Table A.2 (page 85) shows Kendall's τ_B rank correlations

between concept association ratings and the remaining demographic characteristics (see Appendix C for scatterplots of these relationships). Although several of these relationships are significant, most are very small, making them poor candidates for moderators in social priming experiments due to the large sample sizes needed to observe interactions. The three largest correlations are the following: the association between GOD and GENEROSITY increases with political conservatism ($\tau_B = .19$, $z = 7.06$, $p < .001$), the association between HOLDING SOMETHING HOT and INTERPERSONAL WARMTH decreases with age ($\tau_B = -.14$, $z = -5.39$, $p < .001$), and the association between SEXY PICTURES and RELATIONSHIP CLOSENESS decreases with age ($\tau_B = -.13$, $z = -5.10$, $p < .001$).

A.1.3 Discussion

Concept associations studied in social priming research vary greatly in their strength. Moreover, several individual differences emerged as predictors of differences in concept association strength, consistent with the idea that priming effect may differ across different populations (e.g., Cesario et al., 2010). Studies P2 through P5 investigate the hypotheses that (1) more strongly associated concept pairs will be better candidates for priming effects than will more weakly associated concept pairs (Studies P2, P3, and P5) and (2) priming effects will be stronger for those who associate the relevant concepts more strongly (Study P4).

A.2 Study P2

Study P2 tested the priming effect of the strongest concept association pair from Study P1, POWER and BEING IN CHARGE, which had a median association score of 97 out of 100. Participants were randomly assigned to one of two common priming manipulation methods. Half of participants completed a sentence unscrambling task (Srull & Wyer, 1979) and half completed a guided recall manipulation like that used in Studies 2 and 3. This design, combined with a large sample size, ensured the greatest chance to observe a priming effect regardless of which method

is better for obtaining it. After the priming manipulation, participants were asked about their feelings of status, authority, and autonomy to assess how in charge they felt.

A.2.1 Method

Participants. In Study P2, I aimed to collect data from 1400 participants after exclusions. According to simulations, this provides 80% power to detect an interaction where one priming method produced no effect ($d = 0$) and the other produced an effect of size $d = 0.30$. To meet this goal, I recruited 1521 Mechanical Turk workers, who were paid $0.75 each for participation. Of these, 51 workers were excluded for non-US IP addresses (24 workers) or failing an instructional manipulation check (27 workers), leaving 1470 workers (865 women; 605 men) included in analyses. They averaged 30.53 years old ($SD = 10.58$). Most identified as European American (67%). Those with no religious affiliation (47%) and Christians of various denominations (45%) made up the majority of the sample. Most participants (92%) had completed at least some college, and 57% reported an annual income of $35,000 or less.

Procedure. Participants completed a power priming task, dependent measures designed to test their feelings of being in charge, and the same demographics questionnaire as in Study 1.

Priming task. Participants were randomly assigned to one of four conditions in a 2 (Priming Method: Sentence Unscrambling, Recall) x 2 (Priming Condition: Low Power, High Power) design. In the Sentence Unscrambling condition, participants completed a 16-item sentence unscrambling task (Smith & Trope, 2006), in which they had to construct a complete sentence that used four of a list of five words in each item. Eight of the items contained prime words related to low (e.g., *obey*, *servant*, *submits*) or high (e.g., *authority*, *commands*, *controls*) power, depending on participants' assignment to the Low or High Power condition. In the Recall condition, participants completed the same Low or High Power visualization and essay task used in Studies 2 and 3.

Dependent measures. Participants completed 12 items assessing their feelings of being in charge of their lives and their relationships with others. Five questions asked about participants' feelings of status in their personal relationships rated on seven-point Likert scales ranging from *strongly disagree* to *strongly agree* (e.g., "I have a high level of respect in others' eyes"; Anderson, Srivastava, Beer, Spataro, & Chatman, 2006). Four questions asked participants how satisfied they were with their authority when working with others on seven-point scales ranging from *not at all satisfied* to *completely satisfied* (e.g., "The authority I have when I work with others"; adapted from Porter, 1961). Finally, three questions assessed participants' feelings of autonomy from others' opinions, rated on seven-point Likert scales ranging from *strongly disagree* to *strongly agree* (e.g., "I judge myself by what I think is important, not by the values of what others think is important"; Ryff & Keyes, 1995). The orders of the three sets of questions and of items within each set were randomized across participants.

A.2.2 Results

The twelve dependent measures formed a reliable composite index ($\alpha = .84$), so they were averaged to form a single measure of feelings of being in charge. This composite was standardized and regressed on Priming Condition (-.5 = Low Power; .5 = High Power), Priming Method (-.5 = Sentence Unscrambling, .5 = Recall), and their interaction. A main effect of Priming Condition would suggest that priming participants with POWER activates thoughts of BEING IN CHARGE (Low Power vs. High Power). However, an interaction between Priming Condition and Priming Method would indicate that the effect of high vs. low power priming depended on the priming method used. The results of this model are displayed in Table A.3 (page 86). The Priming Condition x Priming Method interaction was significant, suggesting that the effect of the power manipulation depended on the priming method. As illustrated in Figure A.2 (page 91), participants asked to recall a high power situation expressed greater feelings of

being charge (M = 5.08, SD = 0.72) than those asked to call recall a low power situation (M = 4.89, SD = 0. 80; t(688.65) = 3.32, p < .001, d = .25 [.10, .40]), but there was no such difference for between participants who unscrambled sentences containing high (M = 4.89, SD = 0.78) or low power words (M = 4.88, SD = 0.76; t(776.01) = -0.18, p = .855, d = 0.01 [-.13, .15]).

A.2.3 Discussion

In Study P2, priming was observed for a strongly associated concept pair but only for one of two methods. Participants primed with POWER via a recall task reported stronger feelings of BEING IN CHARGE. These results are consistent with the argument that strong concept associations are critical to priming and suggest that strong means of activating these concepts may be important to observing social priming effects. However, testing the claim that strong rather than weak associations yield social priming effects requires also examining a weakly associated concept pair, which is the subject of Study P3.

A.3 Study P3

Study P2 tested the priming effect of a weakly associated concept pair from Study P1, TIME and SELF-REFLECTION, which had a median association score of 21 out of 100. The methods for Study P3 were similar to those of Study P2: Participants were asked to complete either a sentence unscrambling or a recall task to prime TIME or a control concept (MONEY) and then to complete three measures of self-reflection.

A.3.1 Method

Participants. As in Study P2, I aimed to collect data from 1400 participants after exclusions. To meet this goal, I recruited 1513 Mechanical Turk workers, who were paid $0.95 each for participation. Of these, 27 workers were excluded for non-US IP addresses and 39 were excluded for failing an instructional manipulation check (including 3 excluded for both reasons), leaving 1450 workers (744 women; 706 men) included in analyses. They averaged 31.91 years

old (*SD* = 10.84). Most identified as European American (73%). Those with no religious affiliation (48%) and Christians of various denominations (44%) made up the majority of the sample. Most participants (90%) had completed at least some college, and 57% reported an annual income of $35, 000 or less.

Procedure. After consenting, participants completed a priming task, dependent measures to assess their feelings of self-reflection, and the demographics questionnaire used in Study 1.

Priming task. Participants were randomly assigned to one of four conditions in a 2 (Priming Condition: Money, Time) x 2 (Priming Method: Sentence Unscrambling, Recall) design. Participants assigned to the Sentence Unscrambling condition were asked to complete a 17-item sentence unscrambling task (Gino & Mogilner, 2014). Each item in the task consisted of four words, and participants had to form sentences using three of the four words. Ten of these items contained words related to either money (Money condition; e.g., *wallet*, *cash*, *pricey*) or time (Time condition; e.g., *hours*, *minutes*, *clock*). Participants assigned to the Recall condition were asked to asked to recall a situation in which they had a lot of money they could spend as they wished (Money condition) or in which they had a lot of time they could spend as they wished (Time condition). They were guided to vividly relive this experience during a 60 second visualization task and then to write a description of what that had visualized.

Dependent measures. After the priming task, participants completed dependent measures intended to assess their feelings self-reflection. They answered three private self-awareness items from the Situational Self-Awareness Scale (Govern & Marsch, 2001; e.g., "Right now, I am conscious of my inner feelings"), three self-reflection questions developed by Gino and Mogilner (2014; e.g., "Right now, I feel like reflecting on my own life"), and three items assessing the need for self-reflection from the Self-Reflection and Insight Scale (e.g., "It is important to me to try to understand what my feelings mean"; Grant, Franklin, & Langford, 2002). All items were

rated on 7-point Likert scales ranging from *strongly disagree* to *strongly agree*. The order of the questionnaires and of the items within questionnaires was randomized across participants.

A.3.2 Results

The nine dependent measures formed a reliable composite index ($\alpha = .90$), so participants' feelings of self-reflection were scored as an average of these items. This measure was standardized and regressed on effect coded Priming Condition (-.5 = Money, .5 = Time), Priming Method (-.5 = Implicit, .5 = Explicit), and their interaction. Similar to Study P2, a main effect of Priming Condition would suggest that TIME has a social priming influence of SELF-RELFECTION related judgments, and a Priming Condition by Priming Method interaction would indicate that one priming method produced a stronger effect than the other. However, as shown in Table A.4 (page 86), neither the main effect of Priming Condition nor the Priming Condition x Priming Method interaction were significant, and the Priming Condition coefficient was estimated to be approximately 0, suggesting that priming condition did not influence feelings of self-reflection.[6] Among participants who unscrambled sentences containing prime words , the means for self-awareness in the Money ($M = 4.94$, $SD = 1.06$) and Time ($M = 5.01$, $SD = 1.02$) conditions were in the predicted direction for a priming effect, though their difference was nearly 0 and non-significant, $t(755.90) = 0.83$, $p = .408$, $d = 0.06$ [-0.08, 0.20]. However, among participants who completed the recall task, the means for the Money ($M = 5.44$, $SD = 0.89$) and Time ($M = 5.36$, $SD = 0.91$) conditions were in the *opposite* direction, and their difference was larger though still non-significant, $t(678.95) = 1.13$, $p = .261$, $d = -0.09$ [-0.24, 0.06]. This is illustrated in Figure A.3 (page 91).

[6] The effect of Priming Method *was* significant and suggests, unsurprisingly, that asking people to vividly recall and write about events from their lives resulted in higher self-awareness than did unscrambling sentences.

A.3.3 Discussion

Study P3 examined a priming effect when the prime and target concepts are only weakly associated in the sampled population. Participants in Study P1 reported a non-zero but weak association between TIME and SELF-REFLECTION. However, participants in Study 4 who were primed with TIME were no more self-reflective than participants who were primed with MONEY, regardless of the priming method used. These results are consistent with the argument that priming effects will not occur when the prime and target concepts are only weakly related.

A.4 Study P4

Studies P2 and P3 showed that *between* concept pairs, the strength of their association predicts the presence or absence of social priming effects. Study P4 addressed a related question: *Within* a concept association pair, will people who hold the association more strongly show stronger priming effects than those who hold the association more weakly? To answer this question, Study 3 focused on a concept pair from Study 1, GOD and WATCHING, with a moderate association strength (*median* = 41), a wide confidence interval (*CI width* = 9), and a clear demographic predictor of association strength (religiosity). Whereas religiously affiliated participants expressed a moderate association between GOD and WATCHING (*median* = 48), unaffiliated participants saw these concepts as more weakly related (*median* = 37). Half of the participants in Study 3 were primed with God and half with a neutral concept. All were then asked about their self-awareness and self-conscious emotions to assess whether they felt watched. A scale measure of religiosity was included to assess the proposed moderator.

A.4.1 Method

Participants. I aimed to collect data from 600 participants after exclusions. This would provide about 70% power to detect an interaction between religious affiliation (dichotomous) and priming condition, if the priming effect were $d = 0.44$ for religious participants and $d = 0.04$

for non-religious participants, as found in a recent meta-analysis (Shariff, Willard, Anderson, & Norenzayan, 2016). Using a continuous measure of religiosity should considerably increase power to detect this interaction (Royston, Altman, & Sauerbrei, 2006). To meet this goal, I recruited 656 Mechanical Turk workers, who were paid $1.00 each for participation. Of these, 8 workers were excluded for non-US IP addresses and 14 were excluded for failing an instructional manipulation check, leaving 634 workers (408 women; 226 men) included in analyses. They averaged 33.32 years old (*SD* = 11.94). Most identified as European American (69%). Those with no religious affiliation (47%) and Christians of various denominations (43%) made up the majority of the sample. Most participants (93%) had completed at least some college, and 52% reported an annual income of $35,000 or less.

Procedure. Participants completed a priming task, dependent measures to assess their feelings of being watched, a measure of religiosity, and the same demographics questionnaire used in Study 1.

Priming task. Participants were randomly assigned to one of two priming conditions taking the form of a visualization and writing task similar to the Recall priming method from Study 2. Participants assigned to the God condition were asked to reflect deeply on their religious beliefs (or lack thereof), and participants in the Control condition were asked to think about the items in their homes (adapted from McCullough, Carter, Dewall, & Corrales, 2012). Only this priming method was used in Study P4 because of the multiplicative scaling of sample sizes needed to observe higher order interactions (i.e., a three-way Prime x Religiosity x Method interaction) and the ineffectiveness of the Sentence Unscrambling priming method in Study P2.

Dependent measures. Participants answered six questions from the Situational Self-Awareness Scale (Govern & Marsch, 2001), half of which assessed their public self-awareness (e.g., “Right now, I am concerned about the way I present myself”) and half their private self-

awareness (e.g., "Right now, I am conscious of my inner feelings"). These items were rated on seven-point Likert scales ranging from *strongly disagree* to *strongly agree*. Participants also rated how much they felt each of three negative self-conscious emotions (*guilty*, *ashamed*, *embarrassed*) and one positive self-conscious emotion (*pride*) on scales ranging from *not at all* to *extremely*. To distinguish whether any observed effects were specific to self-conscious emotions, as opposed to positive or negative affect more generally, participants were also asked to report how much they felt three negative emotions (*angry*, *afraid*, *disgusted*) and three positive emotions (*joyful*, *surprised*, *interested*). The order of the questionnaires and of the items within questionnaires was randomized across participants.

Religiosity. Participants answered six items about their intrinsic religiosity (e.g., "My personal religious beliefs are very important to me"; Shariff, Cohen, & Norenzayan, 2008) plus a single-item measure of religious attendance ("How often do you attend religious services?"). These items were rated on 7-point scales, with 1 always indicating the least religiosity. A *not applicable* option was also included for several items (e.g., "My religion or faith is an important part of my identity") and scored as 1. Using the same coding scheme for religious affiliation as in Study 1, religious participants ($M = 4.94$, $SD = 1.50$) reported greater religiosity than non-religious participants ($M = 1.85$, $SD = 1.01$, $t(582.81) = 30.53$, $p < .001$, $d = 2.39$, 95% *CI* [2.18, 2.59]).[7]

A.4.2 Results

Mean scores were calculated for each participant on Public Self-Awareness (α = .82), Private Self-Awareness (α = .77), Negative Self-Conscious Emotions (α = .88), Negative Emotions (α = .82), Positive Emotions (α = .62), and the seven Religiosity items (α = .95). Scores for the dependent measures (Public Self-Awareness, Private Self-Awareness, Negative

[7] Four participants refused to answer the religious affiliation question and are not included in this analysis.

Self-Conscious Emotions, and Pride) were only weakly correlated (-.07 < *r*s < .32; see Table A.5, page 86) and so were analyzed separately rather than combined into composite measures.

The standardized dependent variables were regressed on effect coded Condition (-.5 = Control, .5 = God), standardized scale Religiosity, and their interaction. If individual differences in religiosity affect the strength of the association between GOD and WATCHING and thereby the strength of the priming effect, there should be a Condition x Religiosity interaction such that the difference between the God and Control conditions increases with religiosity.

Results of these regressions are shown in Table A.6 (page 87) and illustrated in Figure A.4 (page 92). The predicted Condition x Religiosity interaction was significant only for private self-awareness. A simple slopes analysis at one standard deviation above and below the mean of Religiosity revealed a significant effect of Condition for highly religious participants (b = .61, 95% CI = [.39, .82], $t(630) = 5.58$, $p < .001$) and a smaller significant effect of Condition for weakly religious participants (b = .23, 95% CI = [.02, .44], $t(630) = 2.11$, $p = .035$), suggesting that participants primed with God were more concerned about observation of their thoughts and feelings than were participants in the Control condition and that this effect increased with religiosity. However, a simple Bonferroni correction applied to the four tests of the Priming Condition x Religiosity interaction (reducing α to .0125) makes even this interaction non-significant.

A.4.3 Discussion

The moderately associated concept pair GOD and WATCHING produced a significant (uncorrected) priming effect for one of the dependent measures, and this effect was moderated by religiosity as predicted from the concept associations observed in Study 1. That God priming influenced only private self-awareness is consistent with arguments that people are concerned about supernatural agents' observation of their thoughts, intentions, and motives (Bering &

Johnson, 2005; Johnson & Bering, 2006) but conflicts with previous research finding that God priming increased participants' public but not private self-awareness (Gervais & Norenzayan, 2012). Moreover, that God priming did not increase negative self-conscious emotions (and, in fact, if anything *reduced* them) contradicts a supernatural surveillance interpretation of the results for private self-awareness, in which participants' increased self-awareness is due to concern over the exposure of immoral thoughts. Instead, it seems likely that more religious participants see religious thought as more important or sacred and therefore pay greater attention to it, increasing their private self-awareness when asked to reflect on their religious beliefs (c.f., Kitchens, 2015).

Given this, Study P4 finds ambiguous evidence for social priming moderated by demographic factors related to the strength of the underlying concept association. One possible explanation for this ambiguity is that the association between GOD and WATCHING was not strong enough to produce a social priming effect, even among religious people. Study P5 attempts to address the question of priming for moderately associated concept pairs by testing whether social priming effect sizes increase linearly with association strength or occur only once associations pass a specific threshold.

A.5 Study P5

Like Study 2, Study P5 tested simultaneously the predictability of social priming effects for weakly, moderately, and strongly associated concept pairs selected from the POWER associations in Study P1. Participants were primed with high or low power and made judgments relating to a strongly related concept (BEING IN CHARGE; *median association in Study P1* = 96), a moderately related concept (SELF-ESTEEM; *median* = 47), and a weakly related concept

(DISTRUST; *median* = 21).[8] If social priming depends on concept associations, I should observe an effect of power priming on judgments related to BEING IN CHARGE (as in Study P2) but no effect on judgments related to DISTRUST (as in Studies 2 and 3). Likewise, if a linear model holds for the relationship between concept association strength and priming effect sizes, then I should observe a significant interaction between a linear contrast coefficient and priming condition. But if social priming effects occur only for concept associations that exceed a certain threshold, then I should observe a significant effect only for the strongly related concept pair and interaction with a threshold contrast similar to that used in Study 2.

A.5.1 Method

Participants. As in Study 2, I aimed to collect data from at least 500 participants after exclusions for non-US IP addresses and failing an instructional manipulation check. To meet this goal, I recruited 549 Mechanical Turk workers, who were paid $1.50 each for completing the survey. 13 participants were excluded for completing the survey from non-US IP addresses and 26 were excluded for failing an instructional manipulation check (including one participant excluded for both reasons), leaving 521 participants (350 women; 170 men; 1 not providing gender) included in analyses. They averaged 32.46 years old (*SD* = 10.37). Most identified as European American (74%). Most participants (94%) had completed at least some college, and 45% reported an annual income of $35,000 or less.

Procedure. Participants completed a power priming task; dependent measures designed to assess their feelings of being in charge, their self-esteem, and their distrust of other people; and the same demographics questionnaire used in Study 1.

Priming task. Participants were randomly assigned to one of the same High and Low

[8] In Study P1, BEING IN CHARGE and SELF-ESTEEM were classified as target concepts for POWER (and are plotted in Figure A.1). DISTRUST was included as an abstract unrelated distractor concept.

Power priming manipulations used in Studies 2, 3, and P2. To ensure that participants did not simply copy and paste pre-written power essays, JavaScript was used to disable pasting on the priming task.

Dependent measures. To assess their feelings of being in charge, participants completed the five questions about their feelings of status in their personal relationships and the four questions about their satisfaction with their authority when working with others from Study P2. They also answered four questions about their ability to be authentic in their relationships with others (e.g., "I feel like I can be myself with others") on Likert scales ranging from 1 = *strongly disagree* to 7 = *strongly agree* (Kernis & Goldman, 2006).

Participants' self-esteem was assessed with 10 items. Seven items from the State Self-Esteem Scale asked participants to rate how they felt about themselves at the present moment (e.g., "I feel good about myself"; Heatherton & Polivy, 1991). Three items from the Rosenberg (1965) Self-Esteem Scale asked participants to rate how they feel they compare to others in the present moment (e.g., "I feel that I am at least on an equal plane with others"). And the Single-Item Self-Esteem Scale asked participants directly about their self-esteem: "I have high self-esteem" (Robins, Hendin, & Trzesniewski, 2001). All self-esteem questions were rated on Likert scales ranging from 1 = *strongly disagree* to 7 = *strongly agree*.

Participants also completed three measures of their (dis)trust for others. In addition to the World Values Survey trust measure and the Trust Game used in Studies 2 and 3, participants completed the nine-item Social Values Orientation measure (SVO; Van Lange, Otten, De Bruin, & Joireman, 1997). This measure, like the Trust Game, involved choices about the allocation of raffle tickets between oneself and an anonymous partner. However, in this game, the participant and their partner were each supposedly making simultaneous, independent choices about the division of tickets between them, and both of their choices would affect both of their outcomes.

For each item, participants had to choose one of three ticket allocations. One option, the competitive choice, maximized the difference between the amount given to the participant and the amount given to their partner (e.g., participants gets 480 tickets; partner gets 80 tickets). A second option, the individualistic choice, maximized the amount received by the participant (e.g., participant get 540 tickets; partner gets 280 tickets). A third option, the prosocial choice, maximized the joint outcome—how many tickets were received by both players in total (e.g., participant gets 480 tickets; partner gets 480 tickets). Players receive the largest number of tickets when both partners make the prosocial choice and the smallest number of tickets when they choose the prosocial option but their partners choose the competitive option. Thus, when participants select the prosocial option, they are trusting that their partners will also behave prosocially rather than competitively.

The three sets of dependent measures (feeling in charge, self-esteem, and distrust) were presented in a random order. The orders of the scales used to assess each dependent variable were also randomized, as were the orders of items within each measure.

A.5.2 Results

Full data, including participant demographic information is available upon request. Thirteen participants did not complete the Trust Game and were excluded from analyses involving this measure.

The 13 items measuring feelings of being in charge formed a reliable composite index ($\alpha = .88$) and were averaged to form a single measure. The 10 items measuring self-esteem also formed a reliable composite index ($\alpha = .89$) and were averaged. As in Studies 2 and 3, the three trust measures were scored so that higher numbers indicate greater distrust. For the Social Values Orientation, this was the proportion of items on which participants did *not* choose the prosocial option. The World Values Survey trust measure did not correlate positively with the behavioral

measures of trust, *r*s = -.04, *p*s > .30. However, the number of tickets kept in the Trust Game was positively correlated with scores on the Social Values Orientation scale, $r(506) = .33$, $p < .001$, so these two scores were standardized and averaged to form a composite. The World Values Survey scores were averaged to form a separate composite.

If strong but not weak concept associations produce social priming effects, we should expect to see a significant difference between conditions for the action/inaction goals measure but not for either of the trust measures. Likewise, if priming occurs for moderately related concept pairs, we would expect to see significant differences between conditions on the perspective-taking measures. Means by condition, *t*-tests, and Cohen's *d*s for all dependent measures by condition are reported in Table A.7 (page 88). All *t*-tests use Welch's correction for unequal variances. To control errors via Bonferroni correction, α is set to .0125. The effect of power priming on feelings of being in charge was in the expected direction but smaller in this study than in Study P2 and did not reach significance at the adjusted α level. Comparisons for all of the other dependent variables are close to 0.

As in Study 2, I used three exploratory linear mixed effects models to formally test three predictions about the relationship between association strength and social priming effects.

For the model testing the linear effect of association strength on priming effects, the four dependent measures were assigned a centered and scaled contrast code based on the median observed association strengths in Study P1: -0.26 = POWER and DISTRUST (used for the WVS trust measure and the combined SVO and Trust Game score), 0.01 = POWER and SELF-ESTEEM (used for the self-esteem score), and 0.50 = POWER and BEING IN CHARGE (used for the Action/Inaction Goals scale). Dependent measures were standardized prior to analysis. As shown in Table A.8 (page 88), the fixed effect for the Prime by Association contrast was positive and significant at the $p = .05$ level in this model, though it is not significant when a Bonferroni

correction for the three tests of the interaction is applied ($\alpha = .017$).

In the model testing a low association threshold for priming effects, the dependent variables were coded as 0.50 (being in charge and self-esteem) or -0.50 (the two distrust measures). As shown in Table A.9 (page 88), the fixed effect for the Association by Priming condition interaction was small and non-significant in this model, suggesting that a low threshold for association strength does not represent these data well.

In the model testing a high association threshold for priming effects, the dependent variables were coded as 0.75 (being in charge) or -0.25 (all others). As shown in Table A.10 (page 89), the fixed effect for the Association by Priming condition interaction was positive and significant at $p < .05$ in this model, but a Bonferroni correction reduces this interaction to non-significance, making this only weak evidence for such a threshold.

A.5.3 Discussion

Study P5 tested the hypothesis that strongly, but not weakly, associated concepts would produce priming effects. Although no effects reached significance after accounting for multiple comparisons, the effect sizes are consistent with those in Study 2. Whereas the difference between priming conditions was approximately 0 for most of the associations studied, for the strongest association, between POWER and BEING IN CHARGE, the difference between conditions was only slightly smaller than that seen for the strong associations in Studies 2 and P2. Consistent with this, there is suggestive evidence that a high association threshold best describes the relationship between association strength and social priming effects.

A.6 Summary of Preliminary Results

Studies 1-4 sampled concept pairs from the social psychological priming literature to investigate the hypotheses (1) that priming effects will be found for concept pairs with strong but not weak associations and (2) that priming effects will be moderated by variables that predict

concept association strength. The first hypothesis was supported by the findings that strong (Study P2) but not weak (Study P3) concept associations make good candidates for priming effects. Moreover, evidence from Studies P3 and P4 suggests that associations must be fairly strong for priming effects to occur. The second hypothesis received ambiguous support in Study P3. Future research might settle this question by manipulating concept associations as attempted in Study 3 or by examining moderation where the association is stronger than the one examined in Study P4.

A.7 Tables

Table A.1

Distractor Concepts from Association Task, Study P1

Focal Concept	*Concrete Related*	*Abstract Related*	*Concrete Unrelated*	*Abstract Unrelated*
Power (control and influence over others)	Mayor	Status	Bottle	Distrust
Elderly people	Florida	Old	Water	Organized
Cleanliness	Washing	Spotless	Shell	Aggression
Patriotism	Red	America	Stapler	Talent
God	Church	Lord	Floor	Financial planning
Sexy pictures	Photos	Arousal	Table	Organized
Time	Clock	Hour	Wrapper	Intelligence
Sugary	Strawberry	Dessert	Cement	Competence
Holding something hot	Tea	Heat	Apple	Competence
Death	Funeral	Deceased	Scarf	Organized

Table A.2

Kendall's τ_B rank correlations (and bias-corrected and accelerated confidence intervals based on 10,000 samples) between demographic characteristics and strength of association for concepts pairs from the social priming literature, Study P1.

	Age			*Political Conservatism*			*Subjective SES*			*Objective SES*		
		95% *CI*			95% *CI*			95% *CI*			95% *CI*	
Concept Pair	τ_B	*LL*	*UL*	τ_B	*LL*	*UL*	τ_B	*LL*	*UL*	τ_B	*LL*	*UL*
Power – Being in charge	.01	-.04	.07	-.05	-.11	.00	-.04	-.10	.02	-.04	-.09	.02
Elderly people – Slow walking	-.02	-.08	.03	-.01	-.06	.05	-.03	-.08	.03	-.04	-.09	.02
Cleanliness – White	-.01	-.06	.04	.05	.00	.10	.02	-.03	.07	.03	-.02	.08
Patriotism – Conservatism	**-.08**	**-.13**	**-.03**	**.09**	**.03**	**.14**	.03	-.02	.08	.00	-.06	.05
Power – Self-esteem	-.05	-.10	.00	-.03	-.09	.02	**.08**	**.02**	**.13**	.03	-.02	.07
God – Watching	**-.06**	**-.11**	**-.01**	**.07**	**.02**	**.13**	.04	-.01	.10	-.01	-.06	.04
God – Generosity	.00	-.05	.05	**.19**	**.14**	**.24**	**.09**	**.03**	**.14**	.01	-.04	.06
Sexy pictures – Relationship closeness	**-.13**	**-.18**	**-.08**	-.05	-.10	.01	.04	-.01	.10	-.01	-.06	.04
Time – Self-reflection	.00	-.05	.05	-.02	-.08	.03	.01	-.04	.06	.00	-.05	.05
Sugary – Kindness	.01	-.04	.06	-.01	-.06	.05	.04	-.02	.09	**.05**	**.00**	**.11**
Holding something hot – Interpersonal warmth	**-.14**	**-.19**	**-.09**	**.06**	**.01**	**.11**	**.09**	**.04**	**.15**	-.02	-.07	.04
Death – Selfishness	**-.08**	**-.14**	**-.03**	.00	-.05	-.06	**.07**	**.01**	**.13**	.03	-.03	.08
Death – Generosity	**-.06**	**-.12**	**-.01**	-.02	-.08	.04	.03	-.02	.09	.04	-.02	.09

Note. Bold font indicates $p < .05$. Objective SES is the sum of standardized educational attainment and standardized income.

Table A.3

Linear Model Predicting Feelings of Being in Charge (Standardized) from Priming Condition, Priming Method, and Their Interaction, Study P2

		95% *CI* for *b*			
Predictor	*b*	*LL*	*UL*	*t*	*p*
Priming Condition (-0.5 = Low Power, 0.5 = High Power)	0.13	0.03	0.23	2.52	.012
Priming Method (-0.5 = Sentence Unscrambling, 0.5 = Recall)	0.13	0.03	0.23	2.45	.015
Interaction	0.24	0.03	0.44	2.27	.024
R^2	.01				
F	5.46				

Note. $N = 1470$.

Table A.4

Linear Model Predicting Self-Awareness (Standardized) from Priming Condition, Priming Method, and Their Interaction, Study P3

		95% *CI* for *b*			
Predictor	*b*	*LL*	*UL*	*t*	*p*
Priming Condition (-0.5 = Money, 0.5 = Time)	-0.01	-0.11	0.09	-0.15	.882
Priming Method (-0.5 = Sentence Unscrambling, 0.5 = Recall)	0.43	0.33	0.53	8.29	< .001
Interaction	-0.14	-0.34	0.06	-1.36	.175
R^2	.05				
F	23.68				

Note. $N = 1450$.

Table A.5

Correlations Among Measured Variables, Study P4

Variable	*1*	*2*	*3*	*4*	*5*	*6*
1. Public Self-Awareness						
2. Private Self-Awareness	.16***					
3. Negative Self-Conscious Emotions	.32***	.04				
4. Pride	.00	.18***	-.07			
5. Positive Emotions	.04	.23***	.01	.59***		
6. Negative Emotions	.24***	.06	.71***	-.02	.00	
7. Religiosity	.01	.12***	.01	.27***	.30***	.00

Note. $N = 634$.
*** $p < .001$

Table A.6

Linear Models Predicting Feelings of Being Watched (Standardized) from Priming Condition, Religiosity (Standardized), and Their Interaction, Study P4

		95% *CI* for *b*			
Predictor	*b*	*LL*	*UL*	*t*	*p*
Private Self-Consciousness					
Priming Condition (-0.5 = Control, 0.5 = God)	-0.14	-0.29	0.02	-1.73	.084
Religiosity	0.01	-0.07	0.09	0.20	.846
Interaction	-0.07	-0.23	0.08	-0.89	.372
R^2	.01				
F	1.28				
Public Self-Consciousness					
Priming Condition (-0.5 = Control, 0.5 = God)	0.42	0.27	0.57	5.44	< .001
Religiosity	0.12	0.05	0.20	3.18	.002
Interaction	0.19	0.04	0.34	2.45	.015
R^2	.07				
F	15.26				
Negative Self Conscious Emotions					
Priming Condition (-0.5 = Control, 0.5 = God)	-0.14	-0.29	0.02	-1.75	.082
Religiosity	0.01	-0.07	0.09	0.23	.821
Interaction	0.01	-0.15	0.16	0.07	.942
R^2	.00				
F	1.03				
Pride					
Priming Condition (-0.5 = Control, 0.5 = God)	-0.07	-0.22	0.08	-0.97	.334
Religiosity	0.27	0.19	0.34	6.98	< .001
Interaction	0.14	-0.01	0.29	1.85	.066
R^2	.08				
F	17.64				

Note. $N = 1470$.

Table A.7

Means (Standard Deviations), t-Tests, and Cohen's ds for Dependent Variables, Study 2

									95% *CI*	
Variable	*Low Power*		*High Power*		*t*	*df*	*p*	*d*	*LL*	*UL*
Being in Charge	4.72	(0.91)	4.90	(0.86)	2.32	516.96	.021	0.20	0.03	0.38
Self-Esteem	4.91	(0.98)	4.95	(0.99)	0.49	518.99	.624	0.04	-0.13	0.22
WVS Trust	2.94	(0.43)	2.96	(0.40)	0.39	515.61	.694	0.03	-0.14	0.21
Behavioral Trust	0.00	(0.81)	-0.01	(0.82)	-0.19	505.65	.849	-0.02	-0.19	0.16

Note. Low Power $n = 256$ for Behavioral Trust and 260 for all other variables. High Power $n = 252$ for Behavioral Trust and 261 for all other variables. WVS = World Values Survey. Being in Charge and Self-Esteem are scored on a 1 to 7 scale. WVS Trust is scored on a scale from 1 (most trust) to 4 (least trust). Behavioral Trust is the average of standardized Trust Game and Social Values Orientation scores, with higher scores indicating less trust. *T*-tests use Welch's correction for unequal variances. To correct for multiple comparisons, α is set to .0125

Table A.8

Linear Mixed Effects Model Testing the Prediction of Linear Increase in Social Priming Effects with Concept Association Strength.

		95% *CI* for *b*			
Predictor	*b*	*LL*	*UL*	*t*	*p*
Priming Condition	0.07	-0.04	0.17	1.22	.224
Linear Association Contrast	-0.00	-0.13	0.13	-0.01	.991
Interaction	0.26	0.00	0.51	1.98	.048

Note. $N = 579$. *P*-values are based on the Satterthwaite approximation for degrees of freedom. Confidence intervals are based on profile likelihood approximation.

Table A.9

Linear Mixed Effects Model Testing the Prediction that Social Priming Effects Occur When Concept Associations Exceed a Low Strength Threshold.

		95% *CI* for *b*			
Predictor	*b*	*LL*	*UL*	*t*	*p*
Priming Condition	0.07	-0.04	0.17	1.22	.223
Low Association Threshold Contrast	-0.00	-0.08	0.08	-0.01	.990
Interaction	0.11	-0.04	0.27	1.42	.156

Note. $N = 579$. *P*-values are based on the Satterthwaite approximation for degrees of freedom. Confidence intervals are based on profile likelihood approximation.

Table A.10

Linear Mixed Effects Model Testing the Prediction that Social Priming Effects Occur When Concept Associations Exceed a Moderate Strength Threshold.

		95% *CI* for *b*			
Predictor	*b*	*LL*	*UL*	*t*	*p*
Priming Condition	-0.07	-0.24	0.10	-0.79	.429
High Association Threshold Contrast	-0.00	-0.18	0.18	-0.01	.992
Interaction	0.36	0.00	0.72	1.99	.047

Note. $N = 579$. *P*-values are based on the Satterthwaite approximation for degrees of freedom. Confidence intervals are based on profile likelihood approximation.

A.8 Figures

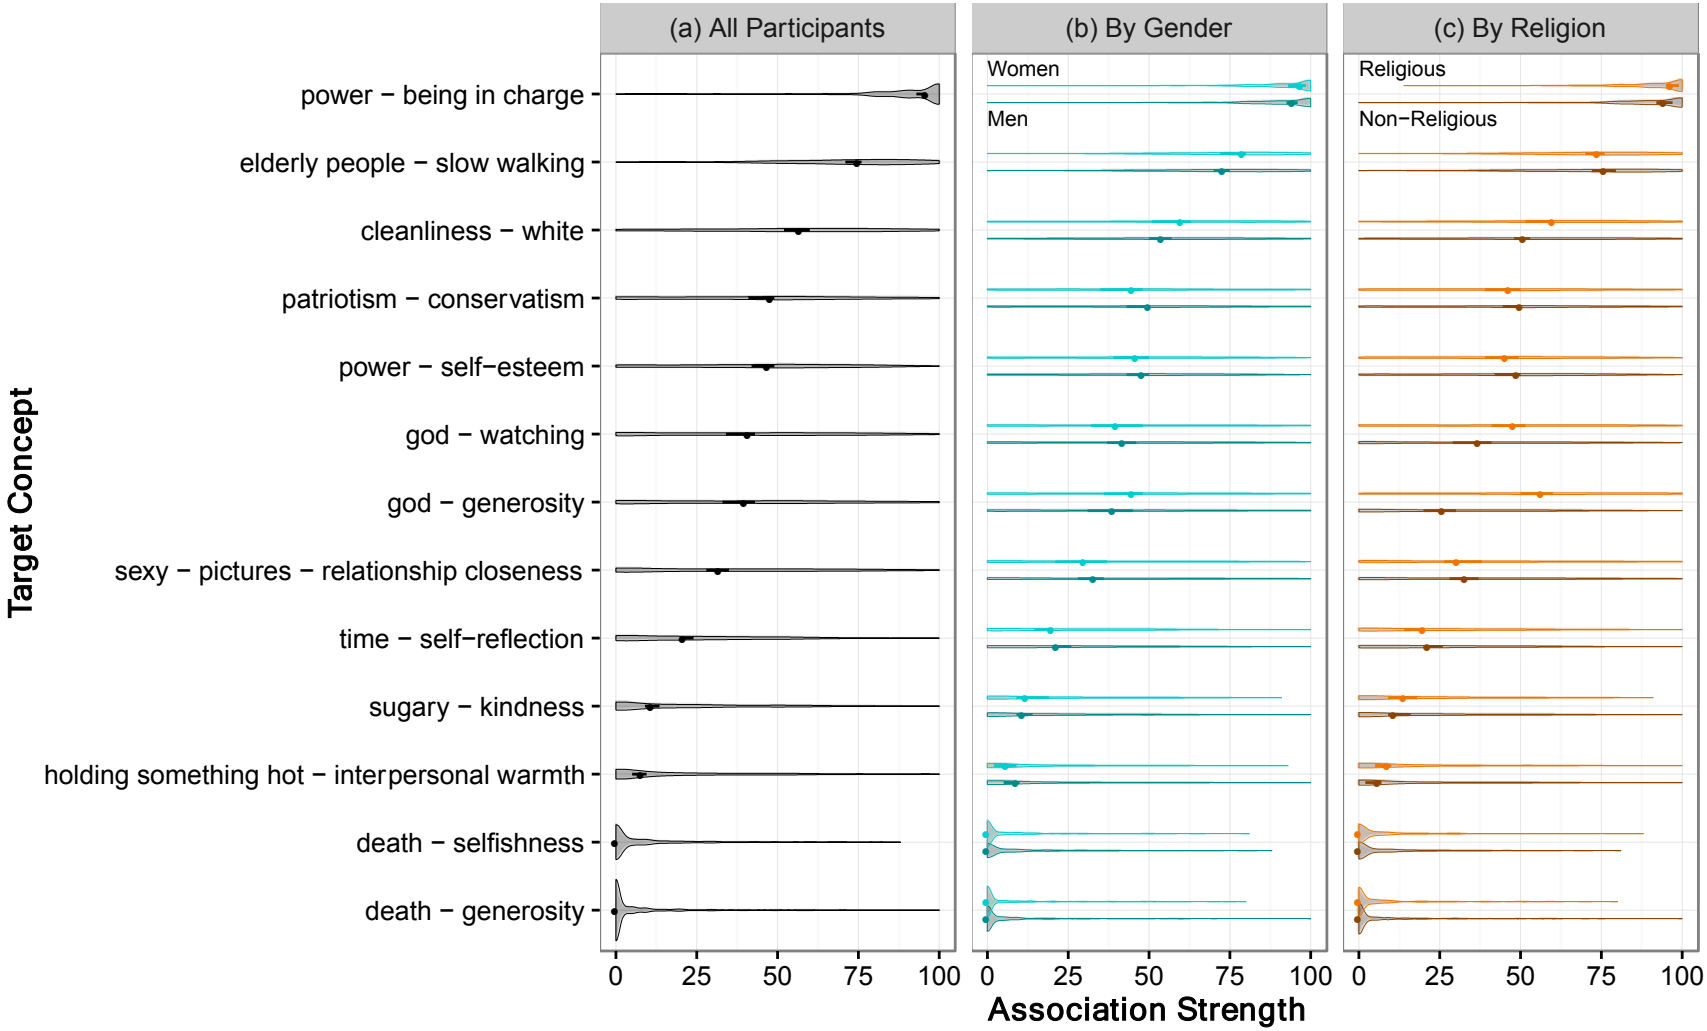

Figure A.1. Median association ratings for 13 concept pairs from the social priming literature for (a) all participants, (b) women and men separately, and (c) religious and non-religious participants separately, Study P1. Error bars represent 95% bias-corrected and accelerated bootstrapped confidence intervals based on 10,000 samples. Shaded areas show kernel smoothed densities.

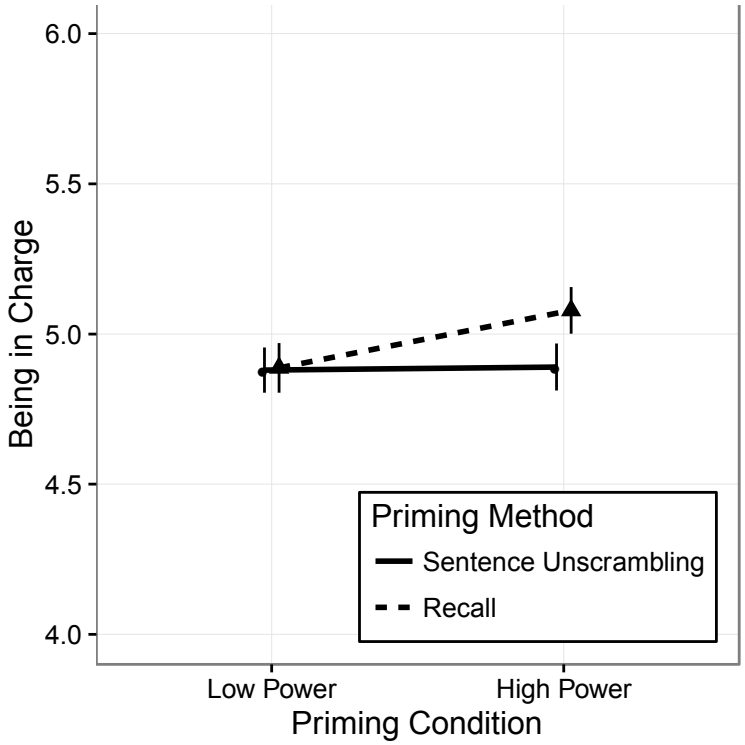

Figure A.2. Observed means of being in charge by Priming Condition and Priming Method, Study P2. Error bars represent 95% confidence intervals.

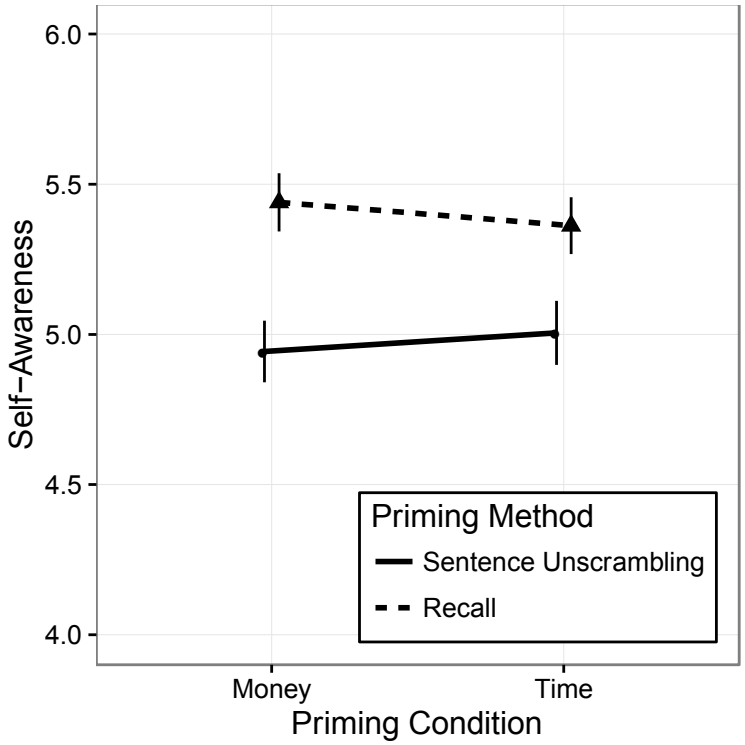

Figure A.3. Observed means of self-awareness by Priming Condition and Priming Method, Study P2. Error bars represent 95% confidence intervals.

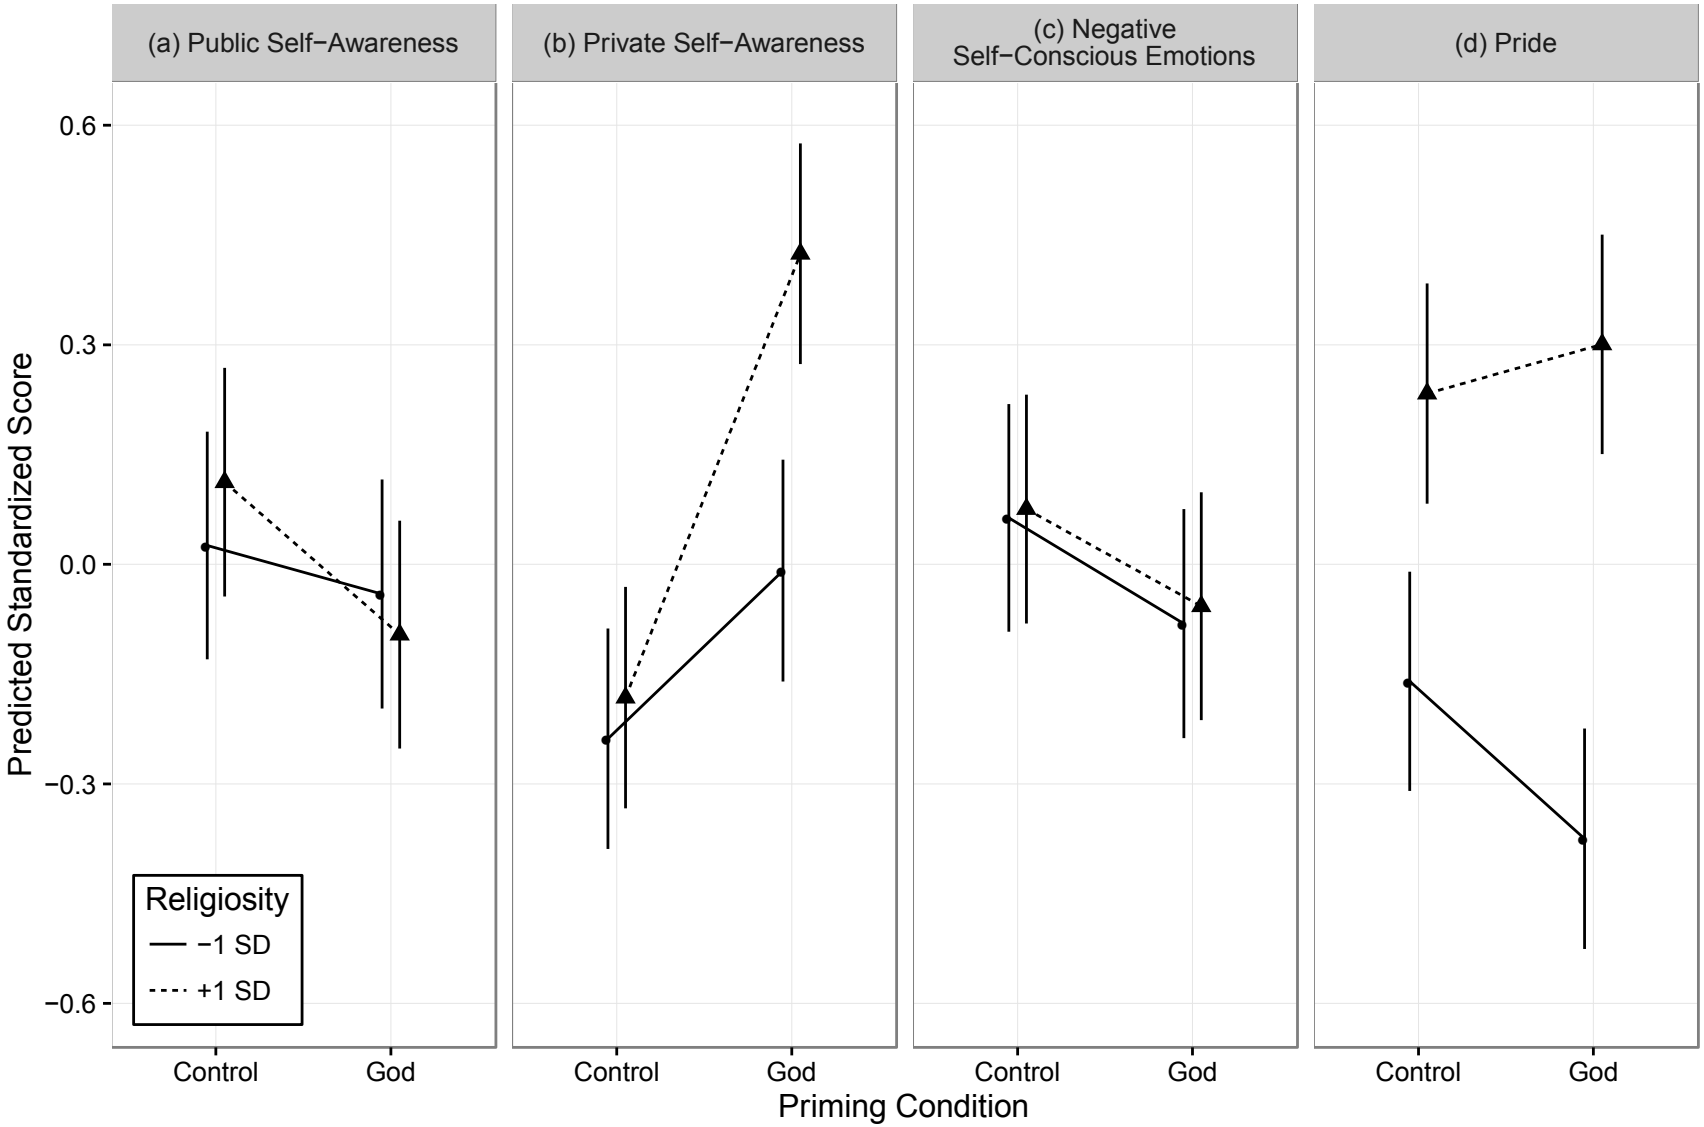

Figure A.4. Predicted standardized values and 95% confidence intervals for four dependent variables by Priming Condition and Religiosity, Study P4.

APPENDIX B: MEANS AND MEDIANS, STUDIES 1 AND P1

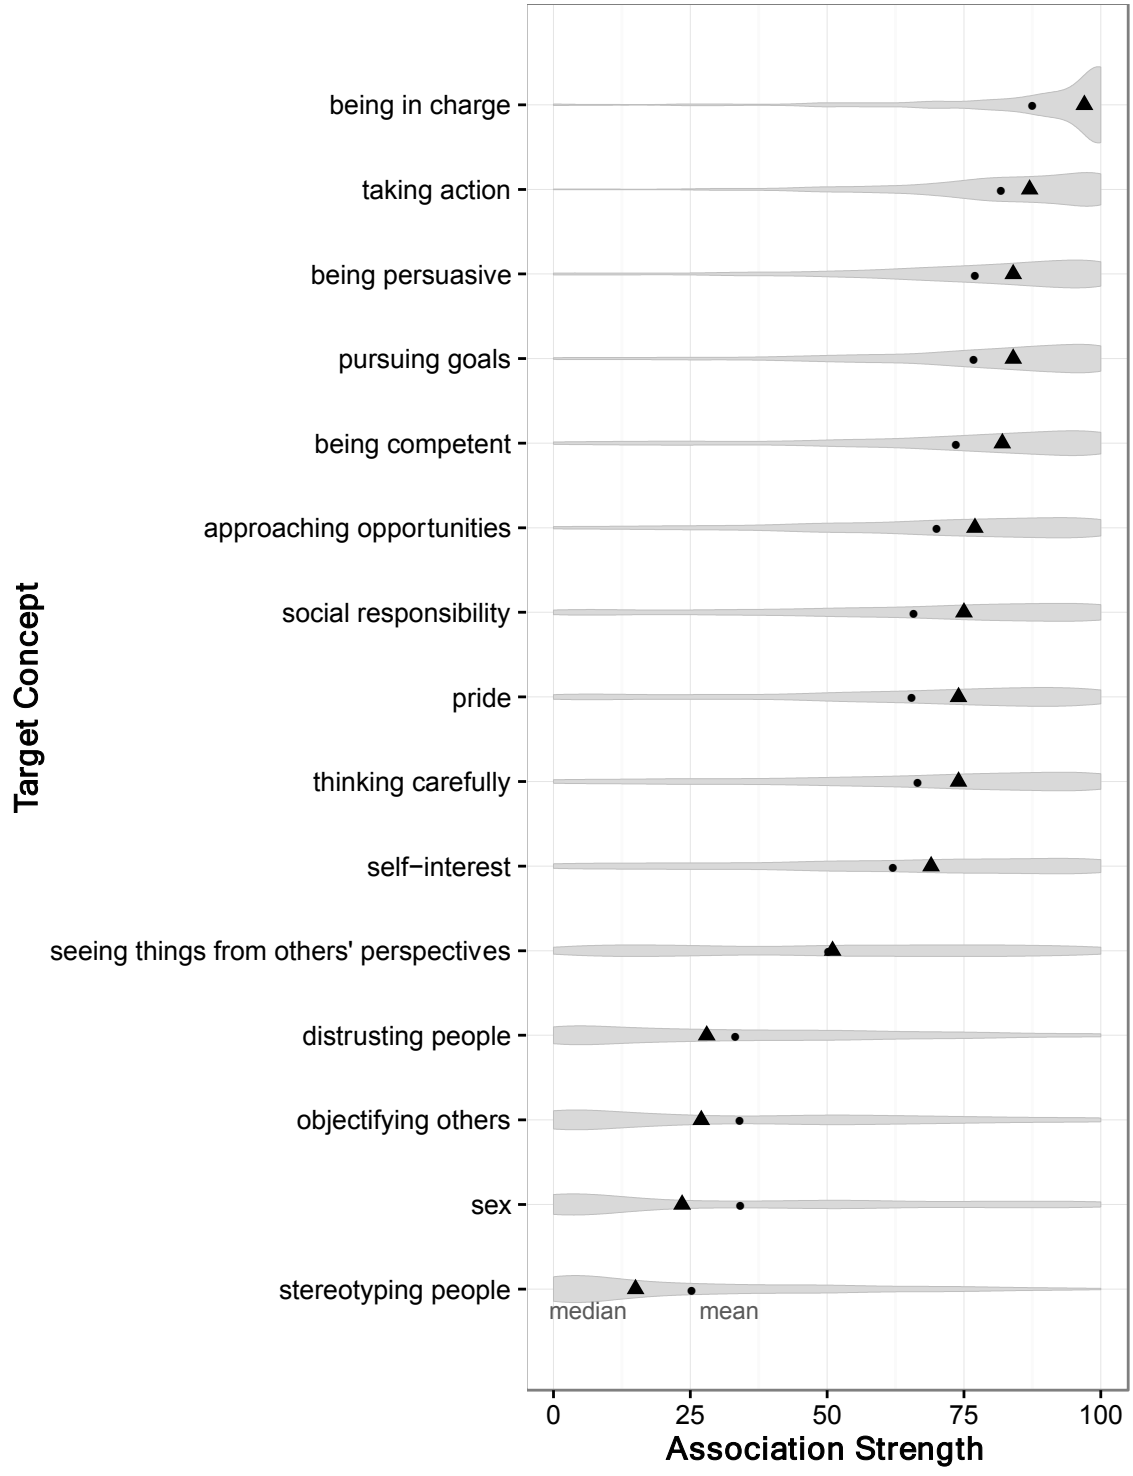

Figure B.1. Mean (circles) and median (triangles) association ratings for 13 concept pairs from the social priming literature, Study 1. Gray regions are violin plots showing the kernel smoothed density for each distribution.

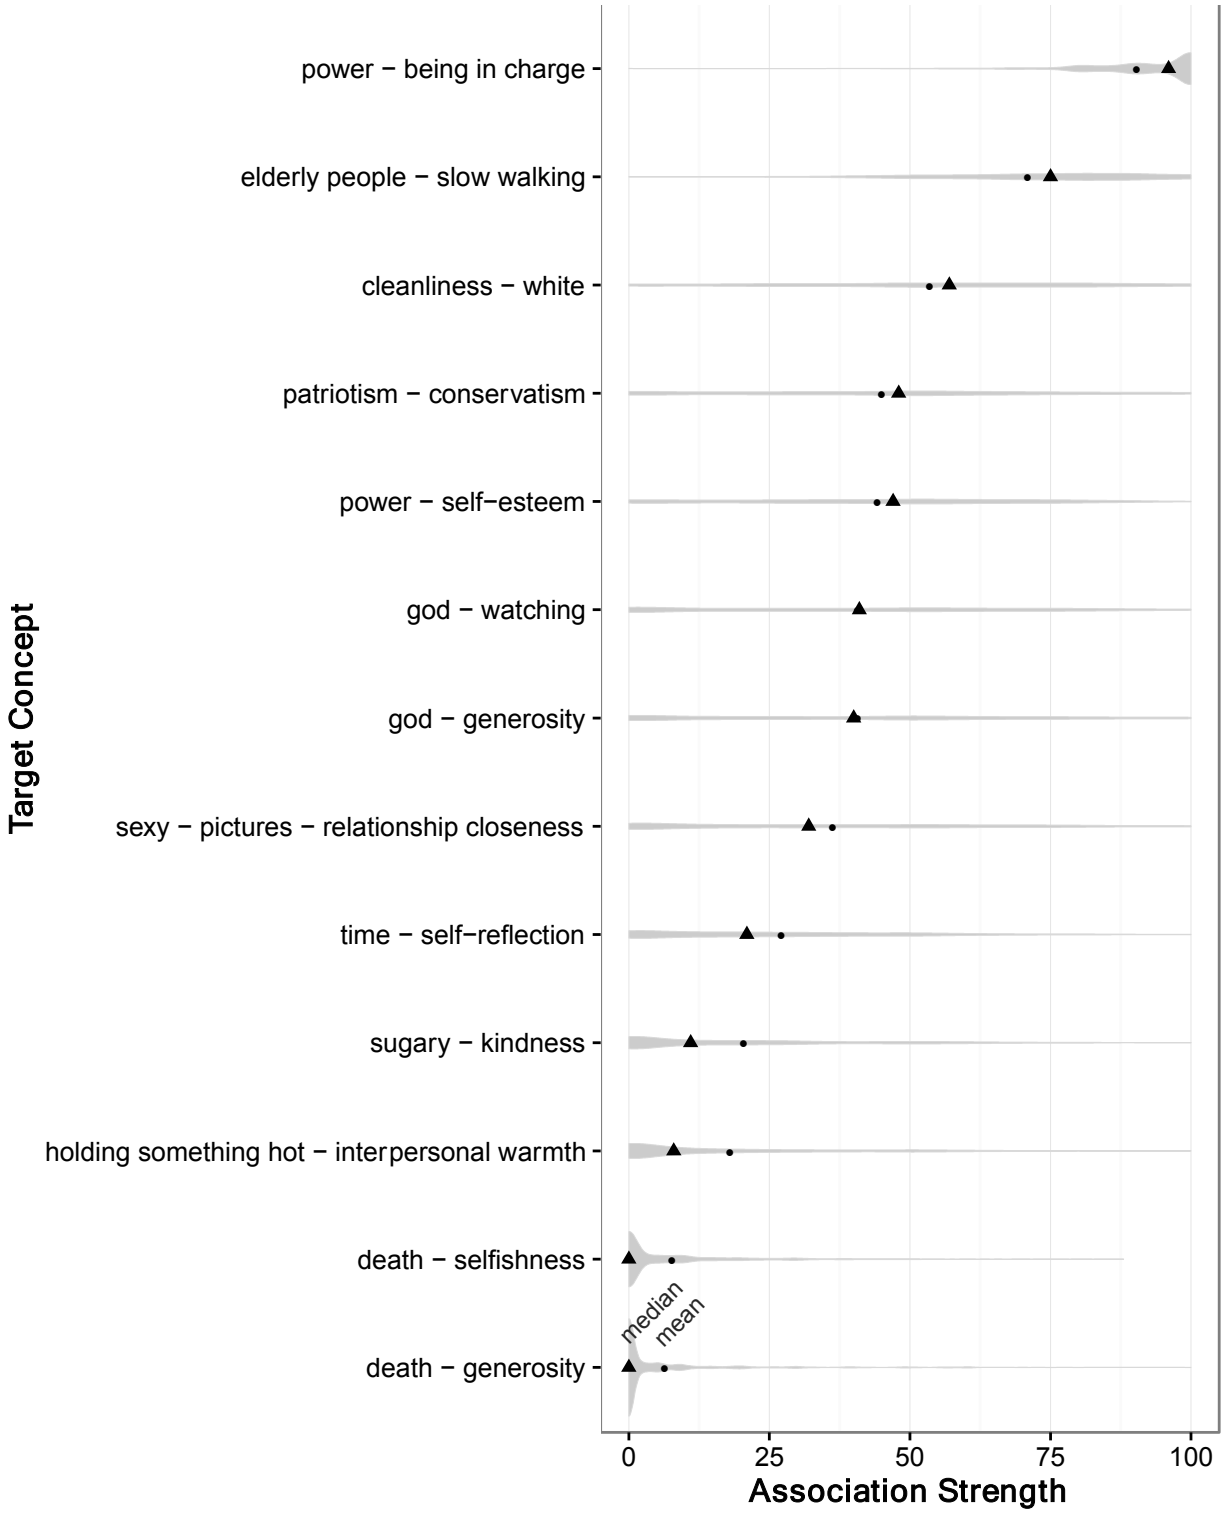

Figure B.2. Mean (circles) and median (triangles) association ratings for 13 concept pairs from the social priming literature, Study P1. Gray regions are violin plots showing the kernel smoothed density for each distribution.