Dissertations, Theses, and Capstone Projects   Graduate Center

5-2018

Perceived Knowledge and Defense of Political Attitudes

Matthew H. Goldberg
The Graduate Center, City University of New York

Recommended Citation

PERCEIVED KNOWLEDGE AND DEFENSE OF POLITICAL ATTITUDES

by

MATTHEW H. GOLDBERG

A dissertation submitted to the Graduate Faculty in Psychology in partial fulfillment of the

requirements for the degree of Doctor of Philosophy, The City University of New York

2018

Perceived Knowledge and Defense of Political Attitudes

by

Matthew H. Goldberg

This manuscript has been read and accepted for the Graduate Faculty in
Psychology in satisfaction of the dissertation requirement for the degree of
Doctor of Philosophy.

_____   _____

Date   Cheryl L. Carmichael, Ph.D.

                                     Chair of Examining Committee

_____   _____

Date   Richard Bodnar, Ph.D.

                                     Executive Officer

Supervisory Committee:

Curtis D. Hardin, Ph.D.

Steven G. Young, Ph.D.

Kimberly Rios, Ph.D.

THE CITY UNIVERSITY OF NEW YORK

Abstract

PERCEIVED KNOWLEDGE AND DEFENSE OF POLITICAL ATTITUDES

by

MATTHEW H. GOLDBERG

Advisors: Cheryl Carmichael & Curtis Hardin

Three experiments tested if perceived knowledge about a political issue predicted people's

willingness to engage with relatively weaker versus stronger belief-threatening information on

that issue. Study 1 assessed people's perceived knowledge on four political issues and, for each

issue, manipulated whether participants chose between arguing against a weak versus moderate

argument or a moderate versus strong argument. Only one issue provided some support for the

primary hypothesis. When people believed they were not knowledgeable about the carbon tax,

giving them a relatively stronger set of arguments to choose from (moderate versus strong)

further increased their preferences for the weaker argument compared to an easier set of choices

(weak versus moderate). When people believed they were high in knowledge of the carbon tax,

the effect disappeared, demonstrating that they were more likely to engage with strong belief-

inconsistent information. Study 2 tested whether perceived knowledge predicted people's

choices of whether to counter-argue versus ignore weak versus strong belief-inconsistent

information. Perceived knowledge predicted preferences for counter-arguing both weak and

strong messages, but the relationship was slightly stronger for strong messages. There was also a

strong effect of condition such that people became significantly more likely to ignore the

information when the information was strong than when it was weak. This effect was primarily

driven by people's expectations that they could not effectively counter-argue strong belief-

inconsistent information. In Study 3, people chose between two debate partners that were either similar in knowledge versus substantially lower in knowledge or between partners that were similar in knowledge versus substantially higher in knowledge. When people believed they knew little about the debate topic (i.e. the death penalty), their choices were unaffected by the available options; when people believed they knew a lot about the topic, they consistently chose the more knowledgeable debate partner. These findings suggest that people's perceived knowledge and their beliefs about their ability to effectively counter-argue belief-inconsistent information play an important role in how they decide to defend their beliefs.

Acknowledgements

My journey to my doctorate started long before graduate admissions. As an undergraduate at Hofstra University, I took a social psychology course with Sarah Novak and quickly learned that this was what I wanted to spend my life studying. There is no one more responsible for my admission to a doctoral program than Sarah. It took endless drafts of my personal statements, getting me the right research experience, and long talks about how much of my "true" self to reveal in the admissions process and graduate life. It turns out that the grit I developed through years of wrestling and mixed martial arts translated to a successful graduate career, where grit was needed most. It was Sarah who celebrated this aspect of me and made sure that I had the opportunities to prove myself in graduate school. Because of Sarah I am writing this, five years later, having earned my doctorate. Thank you so much for believing in me.

There was no one more influential in my graduate career than Cheryl Carmichael. You have gone above and beyond your role as an advisor to teach me what it means to be an excellent scholar, teacher, and mentor. Throughout all of the graduate school growing pains, you brought warmth and kindness—ensuring that I would accomplish every goal I set out to achieve. I will be forever grateful for all that you have done for me and it is no exaggeration when I say I could not have done this without you.

Thank you to Curtis Hardin for shaping my philosophy of science and inspiring so many of my ideas. You have always affirmed that my work was important and worth doing. It is because of you that I learned to think theoretically, giving me the ability to take seemingly disparate findings and integrate them into a unified framework. Your wisdom will always stick with me that the people who are successful in this field are those who do the work because they love doing it.

I couldn't be more grateful for my family for always believing in me. Thank you to my mother, father, and Avi for supporting my every move and for displaying character and work ethic I aspire to match. Thank you to my grandmothers, Bella and Florence, for keeping my life simple. You constantly remind me that all I really need is your love, support, and, of course, too much food.

As one could easily tell from these acknowledgements, I am unbelievably fortunate to have had all of these amazing people shape my life. As I grow, I hope to inspire each of you in all of the unique ways you have inspired me.

Table of Contents

List of Tables

List of Figures

Chapter 1: Introduction

Voter ignorance is a widely recognized phenomenon in American politics (Arnold, 2012; Althaus, 1998). In spite of their ignorance, people still find ways to defend their political beliefs from dissent. Although people can protect their beliefs by expending more effort criticizing belief-inconsistent information (Edwards & Smith, 1996; Ditto & Lopez, 1992) or avoiding it altogether (Eagly & Chaiken, 1993), successfully arguing against strong information provides a greater epistemic benefit than arguing against weak information because it increases attitude certainty (Tormala & Petty, 2002). The proposed research is designed to show that, as people's self-perceived knowledge on an issue increases, their tendency to *choose* to approach strong belief-inconsistent information will increase. Additionally, I propose that people will approach the strongest belief-inconsistent information that they believe they can refute. That is, when people believe they are highly knowledgeable, they will choose to argue against strong rather than weak belief-inconsistent information because refuting strong information proves the superiority of one's beliefs more than refuting weak information.

For example, when people believe they are high in political knowledge they should expect a high probability of successfully counter-arguing criticisms of their political views. Such people should be more inclined to approach strong belief-inconsistent information because counter-arguing strong information increases attitude certainty whereas counter-arguing weak information does not (Tormala & Petty, 2002; also see Tormala & Petty, 2004). People who believe they are low in political knowledge recognize their limited ways of defending their beliefs and should therefore only approach belief-inconsistent information when it is weak—when they know that they can refute it.

Learning about how people defend their political attitudes is crucial to understanding how to effectively relay information to the uninformed electorate. However, the influence of perceived knowledge on information search remains unclear. Thus, it is important to learn more about people's willingness to engage with information that challenges their political position. Cognitive Energetics Theory (CET; Kruglanski, Bélanger, Chen, Köpetz, Pierro, & Mannetti, 2012) provides a guiding theoretical framework for thinking about the kinds of information with which people seek to engage.

Cognitive Energetics Theory

CET uses the early terminology of Lewin (1936) to describe how motivated cognition predicts goal pursuit. That is, it assumes that goal attainment is guided by a *driving force* and a *restraining force*. A driving force represents the maximal amount of energy that a person is willing to invest in attaining a cognitive goal whereas a restraining force limits this energy.

Primary driving forces include the importance of attaining the cognitive goal and the amount of available mental resources. The most prominent example of this is at the core of the Elaboration Likelihood Model (ELM, Petty & Cacioppo, 1986) in which students read arguments in favor of imposing a comprehensive exam for all students at their university or at a distant university. For students in the condition that imposes the exam at their own university, it raises the driving force (i.e. personal importance or vested interest) and thus leads the students to deeply process the message content because the new exam will directly affect them. For students in the condition that imposes the exam at a distant university, there is little vested interest in the outcome, and therefore they will rely on shortcuts in evaluating the information (e.g. source expertise) rather than the argument content. This demonstrates that, when a task is of high importance, people will invest more energy in the task.

The amount of available mental resources is also a critical component of the driving force. For example, Kruglanski and Thompson (1999, Study 3) manipulated whether participants were to evaluate information under high or low cognitive load. The researchers then manipulated whether or not the information came from an expert source (expert vs. inexpert) and the length of the source background information (short vs. long). When under high cognitive load, participants relied more on the brief, but less informative, information to form their judgments whereas participants under low cognitive load relied more on the lengthy, but more informative, information. Inducing high cognitive load reduced participants' driving force towards attaining the goal of accurately judging the relevant information whereas people with the available mental resources (i.e. low cognitive load) were more willing to invest more energy in the task.

On the other hand, a restraining force inhibits goal pursuit via processes such as task difficulty or competing goals. For example, political partisans might be highly motivated (i.e. high driving force) to defend their favorite candidate's tax policy but will be limited in their ability to do so if they face strong counter-arguments from opposing partisans (i.e. high restraining force). Shah and Kruglanski (2002) activated a high driving force in students to perform well on an anagram task by telling them it measured "verbal recall," which the experimenters described as a valuable skill for writing. The researchers demonstrate that, compared to priming task-related goals (i.e. performing well on the anagram task), priming alternative goals (i.e. increased restraining force) led to worse performance on anagram tasks. Activating an alternative goal increased the restraining force on performing well on the anagram task.

Although CET states that the driving and restraining forces are separable, it also states that they are functionally equivalent. That is, for example, raising the restraining force is

analogous to reducing the driving force. CET states that when people are have a high driving force, they seek to choose the most effective way of attaining their cognitive goal even if that means investing more energy. For example, Pierro, De Grada, Mannetti, Livi, and Kruglanski (2004) demonstrate that people rely on cognitive shortcuts (e.g. consensus information) when they have a low driving force but they rely more on argument content when they have a high driving force. Relying primarily on consensus information is fallacious because it does not reveal anything about the arguments themselves. Put simply, people rely on this more efficient, but less valid, technique for understanding information when they are unmotivated to deeply process it (also see Petty & Cacioppo, 1986). On the other hand, when people are highly motivated (i.e. high driving force), they will use more laborious ways of processing information (e.g. assessing message arguments) to the extent that is serves an appropriately important cognitive goal.

Because the primary goal of the current project is to investigate how perceived knowledge predicts people's willingness to engage with information that contradicts their political beliefs, we will select for people that are highly motivated to defend those beliefs. People who are strongly committed (Brehm & Cohen, 1962; Kiesler, 1971) or have publicly advocated for their position (for a review see Harmon-Jones & Harmon-Jones, 2008), for example, will have a strong driving force that motivates them to defend their political beliefs whereas politically unaffiliated people are relatively less likely to be similarly motivated. In the context of this project, the restraining force is the perceived likelihood of successfully refuting the belief-inconsistent information. Thus, people with low perceived knowledge should be less likely to engage with belief-inconsistent information than people with high perceived knowledge because their lower likelihood of successfully counter-arguing it will act like a restraining force.

Perceived Knowledge and Preference for Belief-Inconsistent Information

It is well known that people tend to prefer information that is consistent with their existing beliefs rather than information that is inconsistent with them (Eagly & Chaiken, 1993; for a meta-analysis see Hart et al., 2009). Although the research on selective exposure spans many domains (e.g. personality feedback, majority influence, consumer behavior), political information exposure has been a popular topic in this area. Stroud (2008) found that liberals follow news sources that correspond to liberal positions whereas conservatives follow news sources that correspond to conservative positions. Further, Sweeney and Gruber (1984) demonstrate that Richard Nixon supporters demonstrated less interest and knowledge about the Watergate affair than people who supported a different candidate. People who supported Nixon were motivated to defend their positive view of him and avoiding negative information about him was one way to maintain a favorable attitude towards him.

People demonstrate this preference for belief-consistent information in order to avoid self-threat (Steele, 1988) or psychological discomfort (Festinger, 1957; Brehm & Cohen, 1962; Blanton, Strauts, & Perez, 2012). For example, Steele (1988) explains that people are motivated to maintain a positive view of the self and therefore belief-inconsistent information can be threatening and unpleasant. If a great math student receives a poor grade on a test in which she expected to do well, this elicits self-threat that she may actually not be as good a math student as she believed. The same idea applies to peoples' political positions. If political partisans are confronted with strong negative information about their party's platform, it is threatening because it raises the idea that they are not informed. Thus, she will be motivated to avoid the threat-arousing information.

However, exceptions occur to the general preference for belief-consistent information, particularly when belief-inconsistent information is weak (Lowin, 1967) or if the person is highly

confident that they can defend their views (Albarracín & Mitchell, 2004). For example, Lowin (1967) created and distributed brochures to people during the 1964 presidential campaign that argued in favor of one of the candidates. He manipulated whether the information was consistent or inconsistent with the participant's political attitudes and ease of refutation (weak vs. strong arguments) to determine whether participants would respond to the offer of more information or ignore it. People accepted the offer of more information arguing in favor of their own party's candidate more frequently than information favoring the opposing party's candidate, but only when the arguments were strong. When the arguments were weak, however, people preferred the information arguing in favor of their opposing party's candidate to information favoring their own party's candidate. Put simply, people prefer strong information that supports their beliefs or weak information that goes against their beliefs. People seek out information that goes against their beliefs when they can easily refute it. As opposed to strong belief-inconsistent information, weak belief-inconsistent information has a smaller chance of causing cognitive conflict. Moreover, refuting a weak opposing argument further proves the superiority of one's own argument. Thus, counter-arguing is another way of protecting oneself from the cognitive conflict resulting from belief threat.

Although people may prefer belief-inconsistent information when it is weak as opposed to strong, Tormala and Petty (2002) demonstrated that participants feel more certain after successfully counter-arguing *strong* arguments than weak arguments. This is because defeating strong criticisms reifies one's existing beliefs whereas counter-arguing weak arguments does not. Similarly, counter-arguing against a high-status source increases attitude certainty but counter-arguing a low-status source does not (Tormala & Petty, 2004). People seem inclined to reason, "if my position withstands strong attacks, it must be a good position!" This suggests that counter-

arguing high-quality information provides an epistemic function—assuring people that their committed beliefs remain unchallenged.

As the above evidence suggests, successfully counter-arguing against strong, but not weak, information makes people more certain of their own beliefs. Thus, people who believe they have the necessary knowledge should *choose* to argue against strong belief-inconsistent information because refuting it provides them with an opportunity to increase their certainty. On the other hand, people with low perceived knowledge believe they lack the political knowledge to counter-argue strong information, making weak belief-inconsistent information preferable. This leads us to expect that people who have high perceived political knowledge, and thus a high driving force, would prefer strong rather than weak belief-inconsistent information whereas people with low political knowledge, and thus a low driving force, would prefer weak rather than strong belief-inconsistent information.

The research reported here is novel and innovative in two ways. First, whereas past research has focused on the *consequences* of counter-arguing, this research focuses on *preferences* for counter-arguing. Second, I examined the role of perceived knowledge, which helps us investigate whether people prefer to refute strong belief-inconsistent information *if they believe they have the knowledge to do so*. This sheds light on factors that influence the defensive biases people prefer, such as whether people avoid or engage with information that contradicts their political beliefs.

Overview of the Current Research

In three interrelated studies, I tested how people defended their political attitudes as a function of their perceived political knowledge. In Study 1, I measured perceived knowledge on four issues and manipulated whether people chose between arguing against a weak versus

moderate argument or a moderate versus strong argument. In Study 2, I measured perceived political knowledge and manipulated the strength of the belief-inconsistent information people received (i.e. weak versus strong). Study 2 is novel because the dependent measure tested whether people chose to engage with or ignore the argument (i.e. complete an unrelated task) as opposed to choosing between two arguments. In Study 3, I administered a political knowledge quiz, assigned participants a fake score, and manipulated the ostensible scores of potential disagreeing discussion partners to see if participants chose to engage in a debate with someone who was slightly versus substantially less (or slightly vs. substantially *more*) politically knowledgeable than themselves.

Power Analysis

The novel paradigm of the reported studies made it difficult to determine an informed effect size estimate. According to Cohen (1988), a rule-of-thumb convention for a small odds ratio in the social sciences is 1.49. We used the G*power software to compute the necessary effect size needed to observe this effect with 80% power (for details on G*Power software, see Faul, Erdfelder, Buchner, & Lang, 2009). Assuming an odds ratio of 1.49, we needed a minimum of 214 participants per study.

Chapter 2: Study 1

Study 1 explored people's preferences for arguing against a strong versus relatively weaker argument. This study sought to test if people choose to argue against stronger arguments to the degree that they are knowledgeable on that issue. In order to test the boundary conditions of preferring to argue against stronger arguments (i.e. preferring to counter stronger arguments up to a point), I manipulated whether participants chose between arguing against a weak versus moderate argument or between a moderate versus strong argument. If perceived issue knowledge

predicts preferences for countering stronger arguments, then this study is an appropriate test of how this relationship changes depending on the relative strength of the available arguments.

Method

Participants

Four hundred and fifty-six people (45% female, $M_{age} = 33.45$, $SD_{age} = 9.65$) were recruited through Amazon Mechanical Turk. The TurkPrime software was used to launch the study and restrict access to people in the United States, with a minimal task approval rate of 85%, and a minimum of 100 tasks completed on Mechanical Turk (Litman, Robinson, & Abberbock, 2016). Six participants were removed for failing the attention check (see Materials section below; also see Oppenheimer, Meyvis, & Davidenko, 2009). The final sample contained 450 participants.[1]

Procedure

Participants first answered questions about their positions and perceived political knowledge on 4 policy issues (gun control, death penalty, junk food tax, and carbon tax). Next, participants made a series of 4 forced-choices in which they chose between arguing against a relatively stronger versus weaker message criticizing their political views. For each of the four issues, participants were randomly assigned to choose between a weak and moderate argument or a moderate and strong argument. This was repeated for all four issues. Then participants completed a political knowledge quiz assessing their identification of political public figures as well as their political party knowledge. Lastly, I assessed participants reported behaviors about the frequency of engaging with people who agree or disagree with them on political matters.

Materials

---

[1] Small fluctuations in sample size (2-4 participants) are from analyzing issues separately, in which a given participant had a missing response for one issue but not others. Such differences are reflected in the degrees of freedom reported in each analysis.

Attrition Remedy. After reading a general description of the study, participants completed a task aimed at reducing participant attrition (see Zhou & Fishbach, 2016). Participants read "Many mTurk workers tend to quit once they see the tasks. **If a sizable number of people quit the survey halfway, the data quality of that survey would be compromised. However, our research depends on good quality data.** Thus, before taking this survey, please make sure you do not mind completing the tasks described on the previous page." Then participants were instructed, "Please type below 'I will complete the full survey' so that we can ensure your data will be of high quality. Thank you!" and then there was a text box for them to type the requested sentence and another that asked for their mTurk identification number.

Attention Check. I used an attention check to ensure that participants were not mindlessly responding (Goldberg & Carmichael, 2017). Participants were shown a paragraph of text that appeared to be asking participants to select the political activities in which they regularly engage (e.g. signing petitions, protesting). After three sentences, the message instructed the participant to ignore the instructions and click on an option they would otherwise never choose ("Run for president of the USA") and ignore all other options.

Issue Attitudes. Participants indicated their position on 4 issues (Gun control, Junk food tax, Carbon tax, and Death penalty). For example, "Do you favor or oppose stricter gun control laws?" (*Favor* vs. *Oppose*). For each issue attitude, participants rated their perceived issue knowledge with 4 items (e.g. "How knowledgeable do you feel on this issue?"; "To what extent do you feel you know the most relevant facts about this issue?"; "How much information do you feel you have on this issue?"; "To what extent do you feel you have expertise on this issue?"; $1 =$ *Not at all*, $9 =$ *Extremely*; $\alpha_{\text{Gun control}} = .90$, $\alpha_{\text{Junk food tax}} = .94$, $\alpha_{\text{Carbon tax}} = .95$, $\alpha_{\text{Death penalty}} = .94$; see Rios, Goldberg, & Totton, 2017). Participants also rated the strength of their position ($1 =$

Extremely oppose, 9 = *Extremely favor*). Finally, participants rated the importance of the issue, their attitude correctness, and their attitude clarity (all scales; 1 = *Not at all*, 9 = *Extremely*).

Choice Tasks. Participants were told that they were going make a series of 4 choices between two tasks on each page, and that they would actually complete one of the tasks that they chose. Each choice was between arguing against a relatively stronger and a relatively weaker criticism of the participant's issue position. For each of the four issues, participants were randomly assigned to read a weak and moderate argument or a moderate and strong argument against their position on that issue (for all arguments, see Appendix A).

The survey software automatically displayed the appropriate stimuli (i.e. belief-inconsistent arguments) to participants. Participants were asked to choose which argument (weak vs. moderate; moderate vs. strong) they preferred to argue against. They repeated this choice task for each issue until all four issues were completed. After they completed their last choice, they proceeded to the next screen where the argument they had chosen for the death penalty reappeared along with an essay text box that allowed them to respond to their chosen argument.

The strength of the arguments was determined by pilot testing (N = 75). As in the full study, participants in the pilot were randomly assigned to read a weak and moderate argument or a moderate and strong argument against the participant's position. Weak arguments were defined by the use of only assertions and anecdotes without empirical evidence. Strong arguments were defined by the use of multiple kinds of strong empirical evidence from authoritative sources in which the evidence points to large impacts in the research findings. Moderate arguments were defined by the use of generic sources (e.g. scientists) that pointed to some impact of the findings cited, but less certain and powerful than those cited in the strong arguments. Participants were asked to rate the strength of each argument they received (1 = *Extremely weak*, 7 = *Extremely*

strong). In all cases, a paired samples t-test determined that the moderate argument was rated significantly stronger than the weak argument and the strong argument was rated significantly stronger than the moderate argument (see Tables 1a and 1b).

Political Agreement/Disagreement. Four questions measured people's general exposure to agreeing and disagreeing others on social media and in person. For example, "How often do you [AGREE/DISAGREE] with political opinions and contents that other people post on social media?" (*Never, Only sometimes, Often, Always or nearly always*). Participants also answered the same two questions, one asking about agreement and one about disagreement, regarding their opinion exposure in face-to-face interactions. These items were included for exploratory purposes.

Political Knowledge Quiz. Participants answered 10 factual questions on public figure identification and party knowledge. Questions gauging public figure identification asked about various people in office (e.g. "John Roberts. What job or political office does he now hold?"; "Michael Pence. What job or political office does he now hold?"). Questions gauging party knowledge asked participants to indicate whether a given position is a Democratic or Republican policy position (e.g. "Which party is generally more supportive of restricting abortion?"; "Which party is generally more supportive of a 'path to citizenship'?"). We administered this measure after the forced choices task to ensure that taking the quiz will not affect people's choices about whether to engage with weak, moderate, or strong arguments. Objective knowledge is unlikely to be affected by the choices task and therefore it was preferable to administer the quiz subsequent to the choice tasks.

Demographics. Lastly, participants completed demographic information such as their age, sex, and political ideology (1 = *Very liberal*, 7 = *Very conservative*).

Results

I hypothesized that perceived issue knowledge would predict preferences for arguing against the stronger argument when the choice was between a weak and moderate argument and that this effect would disappear when the choice was between a moderate and strong argument. To test this hypothesis, I entered Condition (0 = Weak vs. Moderate; 1 = Moderate vs. Strong), Issue Perceived Knowledge, and the interaction term into a logistic regression model predicting argument choice (0 = Weaker argument; 1 = Stronger argument). I ran the same model separately for each issue (Gun control, Junk food tax, Carbon tax, and Death penalty).

For the issue of gun control, the Condition × Perceived Knowledge interaction was not significant, (b = -.09, SE = .11), z(442) = -.82, p = .412, 95% CI [-.31, .13]. However, the main effect of Condition was significant, (b = -.40, SE = .20), z(442) = -2.00, p = .045, 95% CI [-.78, -.01]. That is, when participants had a choice between arguing against a weak or moderate argument, participants chose the weaker argument 57% of the time. When they had a choice between a moderate and strong argument, participants preferred the relatively weaker argument (i.e. moderate argument) even more strongly, 67% of the time, OR = 1.49, 95% CI [1.01, 2.19]. The main effect of Perceived Knowledge was not significant, (b = .03, SE = .06), z(442) = .60, p = .550, 95% CI [-.08, .14].

For the issue of the junk food tax, the Condition × Perceived Knowledge interaction was not significant, (b = .10, SE = .09), z(444) = -1.09, p = .275, 95% CI [-28, .08]. The main effect for Condition was not significant, (b = .17, SE = .19), z(444) = .89, p = .372, 95% CI [-.21, .56]. The main effect of Perceived Knowledge was marginal such that higher perceived knowledge was associated with higher odds of choosing to argue against the relatively stronger argument, (b = .08, SE = .05), z(444) = 1.73, p = .084, 95% CI [-.01, .17].

For the issue of the carbon tax, the Condition X Perceived Knowledge was significant, (*b* = .22, *SE* = .09), *z*(446) = 2.44, *p* = .015, 95% CI [.04, .40]. When participants reported low perceived knowledge of carbon taxes (-1SD), the difference between conditions was significant, (*b* = -.85, *SE* = .28), *z*(446) = -2.99, *p* = .003, 95% CI [-1.40, -.29]. That is, when participants received a choice between a moderate and strong argument, it led to 2.34 times the odds of choosing the relatively weaker argument than when participants received a choice between a weak and moderate argument, OR = 2.34, 95% CI [1.34, 4.07][2]. When participants reported high perceived knowledge of carbon taxes (+1SD), the difference between conditions disappeared, (*b* = .12, *SE* = .27), *z*(446) = .43, *p* = .667, 95% CI [-.41, .65]. See Figure 1 for a depiction of the interaction. The main effect of Condition was marginal, (*b* = -.37, *SE* = .19), *z*(446) = -1.89, *p* = .059, 95% CI [-.75, .01]. That is, people preferred the relatively weaker argument 53% of the time when their choice was between a weak and moderate argument but their preference for the weaker argument increased to 62% when their choice was between a moderate and strong argument.

For the issue of the death penalty, the Condition × Perceived Knowledge was not significant, (*b* = .01, *SE* = .10), *z*(446) = -.15, *p* = .879, 95% CI [-.20, .17]. Neither main effect of Condition nor Perceived Knowledge was significant (*p*s > .4)

Discussion

Hypotheses for Study 1 were not confirmed. The main theoretical argument received partial confirmation for only one issue, the carbon tax. When participants reported low perceived knowledge on the carbon tax, their preferences for the relatively weaker argument became significantly stronger when they were given a relatively stronger set of argument choices

---

[2] The signs of the odds ratios reported for this interaction were reversed for ease of explanation and interpretation.

(moderate vs. strong) compared to when they were given a weaker set of choices (weak vs. moderate). This increase in preference for the weaker argument can be interpreted as imposing a greater restraining force on participants, thereby increasing their desire to avoid the stronger argument. Thus, one's desire to avoid the strong argument should be enhanced when the argument is especially strong (in the moderate vs. strong condition) compared to when it is moderately strong (in the weak vs. moderate condition). When participants reported high perceived knowledge on the carbon tax, however, this preference disappeared and leaned towards (although non-significantly) a preference for relatively *stronger* arguments. In effect, participants had a higher driving force to defeat the available arguments and were therefore less influenced by the relatively more difficult set of arguments. Although this does not exactly conform to the original hypothesis, it lends some merit to the idea that people with high perceived knowledge are more likely than people with low perceived knowledge to seek out stronger arguments to argue against.

There was a main effect of condition for both gun control and carbon tax. In both cases the main effect demonstrated that people modestly preferred to argue against a weak argument to a moderate argument but preferred to argue against the moderate argument at a substantially higher rate when the alternative was a strong argument. Simply put, people demonstrated stronger preferences for the relatively weaker argument when they chose between a stronger set of arguments (e.g. moderate vs. strong) compared to when they choose between a weaker set of arguments (e.g. weak vs. moderate). The fact that the effects of condition were modest and not moderated by perceived knowledge in three of the four issues suggests that the above findings be interpreted with caution.

A potential methodological flaw that may explain why this study's hypotheses were not confirmed is that participants' perceived issue knowledge might not adequately capture whether participants believe they could refute the specific information in the belief-inconsistent arguments. For example, it is entirely possible that people with high perceived knowledge prefer to argue against stronger arguments they believe they can refute, but our perceived knowledge measure did not assess whether participants believed they could refute *the specific message they received in the study*. One way to address this issue is to include questions that ask participants about their perceived ability to counter-argue the specific message they receive. If indeed the main flaw of Study 1 was that the perceived knowledge measure did not appropriately measure one's perceived ability to counter-argue to messages they read, then using a better measure in Study 2 would produce more convincing results. For example, it may be better to ask, "how likely are you to effectively counter-argue the argument above?" than the general knowledge questions we asked in Study 1 such as "how much knowledge do you feel you have on this issue?" because the former asks about the specific argument in the task whereas the latter asks about the participants' general perceived knowledge.

Additionally, more research is needed to explore the factors that predict whether someone will use one defense strategy over another strategy (e.g. counter-argue versus ignore) as opposed to arguing against one message over another. Giving participants more flexibility in how to defend their beliefs should give us more insight into what they would do if they were unconstrained by the experimental procedures. This is the main contribution of Study 2.

Chapter 3: Study 2

The purpose of Study 2 was to test how the strength of a belief-inconsistent message influences how people defend against it, such as deciding whether to counter-argue the message

or ignore it. Study 2 explored a situation in which participants rated their perceived issue knowledge, were randomly assigned to receive either a weak or strong argument against their issue position, gauged their ability to counter-argue the argument, and had the choice between counter-arguing the message on the screen or to complete an unrelated task. If people are more interested in counter-arguing against stronger arguments if they have the knowledge to do so, then I hypothesize that perceived knowledge will predict preferences for counter-arguing when participants receive a strong belief-inconsistent message but will predict preferences for ignoring the message when it is weak.

Method

Participants

Four hundred and twenty-eight people participated in this study through Amazon's Mechanical Turk. The TurkPrime software (Litman, Robinson, & Abberbock, 2016) was used to make the same participant restrictions as Study 1 (from the U.S.A., minimum 85% task approval, minimum 100 tasks completed). Ten participants were removed for failing the attention check. The final sample contained 418 participants[3].

Procedure

In Study 2, participants answered the same questions as in Study 1 that tap into their policy positions and their perceived issue knowledge for the issues of gun control and the death penalty. Next, participants were randomly assigned to receive a strong or weak argument against their political position on one of the issues. Participants then rated their perceived likelihood of successfully counter-arguing the message and then chose whether they preferred to provide a

---

[3] Small deviations from this sample size for some analyses are due to participant missing data from some variables but not others. I sought to maximize the sample size by using all available data for that particular analysis.

written argument against the message or complete a writing task unrelated to their political beliefs. The dependent variable was participants' decision to argue versus ignore the message.

Materials

Argument Manipulation. Participants were randomly assigned to receive either a weak or strong argument that criticized their position on either gun control or the death penalty. We reduced this study to two issues because gun control and the death penalty are likely to be more familiar to participants, and therefore participants are more likely to have a clearer position on these issues than on the issues of junk food or carbon taxes. Additionally, given the same sample size, it is more statistically powerful to use two issues instead of four. The strong and weak arguments for these issues were identical to those in Study 1.

Choice Task. Participants were instructed to choose between arguing against the information or completing another unrelated task that has been equated for interest and difficulty. Participants read, "The task on this page is to spend the next 3 minutes arguing against the message below. However, we understand that some participants might want to complete a different task instead. We pretested several other tasks to ensure they take the same amount of time and effort so you may choose to complete the current task or an alternative unrelated task."

Expectations of Effective Counter-Arguing. After reading the argument criticizing their position on the randomly selected issue (Gun control or the Death penalty), participants answered four questions measuring their expectations of effective counter-arguing (see Albarracín & Mitchell, 2004). Participants rated how likely they were to "effectively counter-argue the argument above," "effectively refute the argument above", "identify problems in the argument above", "know counter-arguments to the argument above" (1 = *Extremely unlikely*, 9 =

Extremely likely). These four items were averaged to form a composite measure of expectations of effective counter-arguing ($\alpha = .95$)

Instructions for Chosen Task. If participants chose to complete the counter-arguing task, they were instructed on the next page "Please spend the next three minutes arguing against the message above." If participants chose to complete the unrelated task, they were instructed on the next page "Please spend the next three minutes describing what political issue is most important to you and why." After making their choice, participants completed their chosen task and answered basic demographic questions.

Results

I hypothesized that, when participants receive a strong criticism of their position, perceived issue knowledge would positively predict a preference for arguing against the criticism (vs. choosing to do the unrelated task). When participants receive a weak criticism, I predicted that perceived issue knowledge would negatively predict a preference for arguing against the criticism (i.e. preference for doing an unrelated task). To test this hypothesis, I first tested the Condition (Weak vs. Strong) × Perceived Knowledge × Issue (Gun control vs. Death penalty) interaction to ensure that the hypothesized effects were not moderated by issue. It was not significant, ($b = -.19$, $SE = .28$), $z(410) = -.68$, $p = .499$, 95% CI [-.74, .36] and therefore I collapsed across issue and tested the Condition (Weak vs. Strong) × Perceived Knowledge interaction predicting task choice (0 = Counter-argue task, 1 = Unrelated task). Although the interaction was not significant, ($b = -.19$, $SE = .14$), $z(414) = -1.39$, $p = .163$, 95% CI [-.46, .08], the pattern of results lend partial support to my theoretical argument. Unexpectedly, when participants received a weak argument, perceived knowledge significantly predicted preferences to counter-argue over completing an unrelated task, ($b = -.19$, $SE = .08$), $z(414) = -2.47$, $p = .014$,

95% CI [-.34, .04]. That is, each one unit increase in perceived knowledge led to 21% greater odds of choosing to counter-argue, OR = 1.21, 95% CI [1.04, 1.41]. In line with the hypothesis, however, when participants received a strong argument, this relationship got even stronger, (b = -.38, SE = .11), $z(414)$ = -3.35, $p < .001$, 95% CI [-.61, -.16]. That is, each one unit increase in perceived knowledge led to a 47% greater odds of choosing to counter-argue, OR = 1.47, 95% CI [1.17, 1.84] (See Figure 2).

The main effect of condition was significant such that participants were significantly more likely to choose the unrelated task (i.e. ignore the argument) when they received a strong argument than when they received a weak argument, (b = 1.03, SE = .25), $z(414)$ = 4.14, $p < .001$, 95% CI [.54, 1.52]. That is, when participants received a strong argument it led to 2.80 times the odds of choosing the unrelated task over counter-arguing, OR = 2.80, 95% CI [1.72, 4.57]. The main effect of perceived knowledge was significant such that participants were significantly more likely choose to counter-argue as perceived knowledge increased, (b = -.28, SE = .07), $z(414)$ = -4.16, $p < .001$, 95% CI [-.41, -.15].

In an attempt to more precisely test if people seek to counter-argue messages only to the degree they believe they could do so successfully, I used the PROCESS macro in SPSS (Model 4; Hayes, 2013) to test if the relationship between Condition (X; Weak vs. Strong) and Task Choice (Y; 0 = Counter-argue task, 1 = Unrelated task) was mediated by their perceived expectations of effective counter-arguing (M; continuous). Compared to when participants received a weak argument, giving participants a strong argument increased participants' preferences for the unrelated task indirectly through their perceptions that they could not effectively counter-argue the message. Put another way, giving participants a weak belief-inconsistent message increased their belief that they could effectively counter-argue it, which in

turn predicted preferences for counter-arguing over completing the unrelated task (i.e. ignoring the message; see Figure 3). Bootstrapping revealed that the indirect effect was significantly different from zero ($b = .61$, $SE_{boot} = .14$), 95% CI_{boot} [.38, .93].

Discussion

Study 2 findings confirm the hypothesis that, when participants received a strong argument, perceived knowledge predicted preferences for counter-arguing as opposed to completing an unrelated task. I did not expect to find that, when the argument was weak, perceived knowledge also predicted preferences for counter-arguing. Although this was not expected—and the interaction not quite significant—these results provide some support for my theoretical argument that people prefer to argue against a strong argument to the degree they believe they have the knowledge to do so and that this relationship was somewhat weaker when participants received a weak argument. I proposed that this is because people become more certain after arguing against strong (as opposed to weak) arguments (Tormala & Petty, 2002; 2004). Thus, people should see greater value in counter-arguing a strong criticism of their beliefs than a weak one and that people would choose to counter-argue accordingly. It is also possible, however, that perceived knowledge was a weaker predictor in the weak condition because most people believed they could successfully counter-argue.

It is worth noting that there was no level of perceived knowledge at which participants' majority choice was to counter-argue. In other words, participants always chose to ignore at a higher rate than counter-argue. However, as perceived knowledge increased, participants' preferences to counter-argue also increased.

The analysis of mechanism was promising. Compared to when participants received a strong argument, receiving a weak argument cause participants to increase their expectations that

they could effectively counter-argue which in turn led participants to increase their preference for counter-arguing instead of ignoring. The fact that expectations of effective counter-arguing strongly predicted preferences for counter-arguing suggests that there is epistemic value in counter-arguing over ignoring. If people's cognitive goal is to feel certain about their beliefs, Cognitive Energetics Theory would predict that people will invest the necessary energy to maintain certainty even if the task is more laborious than alternative tasks. Ignoring the information leaves the cognitive discrepancy intact whereas counter-arguing does not. Thus, it is more defensively advantageous to counter-argue than to ignore *if one believes they can counter-argue successfully*. In terms of CET, people who believe they will be successful in counter-arguing have a high driving force and will be more likely to choose the task that fits their cognitive goal of maintaining attitude certainty.

Taken together, these findings are encouraging for future research that will help identify the factors that predict people's choice of belief-defense. Given that people still largely chose to ignore the belief-inconsistent information suggests that information avoidance will continue to be an obstacle in fostering intergroup communication.

One limitation of Study 2 is that participants got to see the argument that they might choose to argue against (versus ignore) before making their decision. Thus, they were likely fully aware of whether or not they could successfully counter-argue the message. When engaging with disagreeing others, it is not always clear what the other person might argue in defending their position, so this might breed more or less perceived "risk" of arguing with the other person. In Study 3, I use a novel paradigm to manipulate the perceived issue knowledge of the participant's potential debate partner, which should provide more insight into the differences between choosing arguments versus choosing someone to debate.

Chapter 4: Study 3

In Study 3 I investigated people's choice of debate partner depending on the partner's ostensible issue knowledge. In Studies 1 and 2 I manipulated participants' exposure to arguments of varying strength. In Study 3, on the other hand, I manipulated participants' choices of a debate partner. I reasoned that people would be less interested in debating someone with substantially lower knowledge than themselves because there would be little to epistemically gain from the debate. If people have a high driving force to maintain or even increase attitude certainty, then this goal would be best achieved by arguing against someone who is relatively high in knowledge (Tormala & Petty, 2002; 2004), but not so high as to pose a risk of decreasing certainty. Thus, I expected that participants would choose a debate partner with similar issue knowledge over one with substantially lower knowledge. However, I expected that when people had a choice between someone with similar issue knowledge or substantially more knowledge, it would create a greater risk of their beliefs being truly challenged if they were to choose the substantially more knowledgeable partner. Therefore, people would again prefer a debate partner with similar knowledge to oneself.

Method

Participants

Four hundred and fifty two people participated in this study through Amazon Mechanical Turk (46% female, M_{age} = 36.22, SD_{age} = 11.94). Like Studies 1 and 2, I used the TurkPrime software (Litman, Robinson, & Abberbock, 2016) to limit study access to people in United States, with a minimum task approval of 85%, and a minimum of 100 tasks completed on Mechanical Turk. Twelve people were removed from analyses because they did not believe the discussion would be real.

Procedure

First, participants answered basic demographic questions such as age and sex. Then

participants indicated their position on gun control and the death penalty. For each issue,

participants answered questions measuring their perceived knowledge for each issue. Then

participants took a quiz that ostensibly measured their knowledge on the death penalty.

Participants then proceeded to the next page that stated that their results were being calculated.

After being given false feedback about their knowledge score (which was identical for all

participants), I manipulated participants' choice between two debate partners, both of whom

always had the opposite opinion of the participant. The two potential debate partners' death

penalty knowledge scores were also displayed. Participants were randomly assigned to choose

between 1) a discussion partner with a similar death penalty knowledge score and one with a

significantly lower score or between 2) a discussion partner with a similar death penalty

knowledge score and one with a significantly higher score. Participants made their choice,

completed a suspicion probe, and then were informed that there would be no debate.

Materials

Demographics. First, participants indicated their age, sex, and political identification (1

= *Very liberal*, 7 = *Very conservative*).

Issue Positions. Then participants answered the same questions from Studies 1 and 2 on

their position on gun control and the death penalty (*Favor* vs. *Oppose*) and four items for each

issue measuring their perceived knowledge ($\alpha = .93$).

Political Knowledge Quiz. Participants took a political knowledge multiple choice quiz

and were given a score that ostensibly measured their political knowledge on the death penalty.

The quiz was constructed using easy, moderate, and difficult questions so that assigning all

participants a moderate score (i.e. 67th percentile) was believable (for all quiz questions and answer choices, see Appendix B). For example, an easy question was "What is the most accurate definition of capital punishment?" whereas a difficult question was "What is the average cost to tax payers of life in prison?" On the following page they read "Your quiz results are being calculated. Please move on to the next page to view your results."

Once participants move on to the next page, they read "The results of this quiz indicate that you answered 6/10 of the questions correctly. This means that you scored **higher** than 67% of your peers." The participant's score was held constant for all participants.

Debate Partner Manipulation. On the following page, participants read that they would be debating the topic of the death penalty with another worker from Mechanical Turk (for a similar procedure, see Rios, Goldberg, & Totton, 2017). They read:

> "In this part of the study, we are interested in how people talk about controversial issues with others. You will participate in an online discussion with another person who is currently logged into this study.
>
> Recently, we have taken advantage of a special new feature in our survey design system, which allows us to connect you with people who are also currently participating in this study. Because this is such a massive study, and online experiments on Mechanical Turk are so convenient and popular, it generally takes less than a minute to connect you via our chat feature to another participant **who has the opposite opinion as you do** on the death penalty.
>
> For this discussion, you can choose between two people that have been randomly selected by our survey system. Please go to the next screen to make this choice."

On the following page, participants were reminded of their knowledge score and told that a choice of two possible debate partners were randomly selected from which they can choose to debate. For both potential debate partners, an 8-digit identification number, the person's position on the death penalty, and the person's quiz score (measuring knowledge of the death penalty) were displayed. The debate partners' scores were shown in a similar format to the participant's

score. For example, "Quiz score: 3/10, which is <u>better than 36% of your peers.</u>" Participants were randomly assigned to one of two sets of partner choices. The choice between discussion partners was either 1) a discussion partner with a similar knowledge score (64^{th} percentile) vs. one with a substantially lower score (36^{th} percentile) or 2) a discussion partner with a similar knowledge score (70^{th} percentile) vs. a substantially higher score (98^{th} percentile). In all cases, participants chose between two discussion partners that both have the opposite opinion on the issue that will be discussed.

 Suspicion Probe. After participants made their choice of discussion partner, they were given the following suspicion probe: "Before we begin, please jot down as specifically as you can what you think will happen during the online discussion." If participants expressed that they did not believe the discussion would be real, they were removed from the sample (see Participants section above). Then participants were told that they would not be having a discussion with another person and were informed about the purpose of the study.

<div align="center">

Results

</div>

 I hypothesized that participants would choose a discussion partner with a similar knowledge score (i.e. close to 67^{th} percentile) when the alternative choice was a partner with substantially lower knowledge. Further, I hypothesized that participants would choose a discussion partner with a similar knowledge score when the alternative choice was a partner with substantially higher knowledge. To test this hypothesis, I entered condition (0 = extremely low knowledge alternative; 1 = extremely high knowledge alternative) into a binary logistic regression predicting choice of partner (0 = similar to own knowledge; 1 = alternative). Further, I added participants' ratings of their perceived knowledge on the death penalty as a covariate to control for individual differences in perceived knowledge that existed before the fake quiz

feedback was given to participants. As expected, participants chose the discussion partner with similar knowledge on the death penalty more frequently than a substantially less knowledgeable alternative. Unexpectedly, however, participants chose the substantially more knowledgeable alternative partner more frequently than the similar knowledge partner. Put simply, participants in both conditions chose the more knowledgeable partner. The main effect was significant, ($b =$ .15, $SE = .05$), $z(437) = 3.11$, $p = .002$, 95% CI [.05, .24]. Going from the condition in which the alternative partner was substantially less knowledgeable to the condition in which the alternative partner was substantially *more* knowledgeable led to 16% greater odds of choosing the alternative, $OR = 1.16$, 95% CI [1.06, 1.27]. That is, people preferred a similarly knowledgeable partner to an alternative that was substantially less knowledgeable but increased their preference for an alternative partner when they were substantially more knowledgeable than the participant.

I also tested the Condition × Perceived Knowledge interaction to test if the effect of condition was weaker or non-existent for participants with low perceived knowledge and stronger for participants with high perceived knowledge. I ran the same binary logistic regression model with the Condition × Perceived Knowledge interaction term. The interaction was marginally significant, ($b = .16$, $SE = .10$), $z(436) = 1.62$, $p = .105$, 95% CI [-.03, .35]. The effect of condition was not significant when participants had low perceived knowledge (-1SD; perceived knowledge score = 3.67 out of 9), ($b = .28$, $SE = .27$), $z(436) = 1.04$, $p = .300$, 95% CI [-.25, .82]. However, the effect of condition got progressively stronger as perceived knowledge went up and became significant once participants reached 4.61 out of 9 on the perceived knowledge scale, ($b = .43$, $SE = .22$), $z(436) = 1.96$, $p = .050$, 95% CI [.00, .86]. This effect remained significant and continued to get stronger for participants as they reported higher perceived knowledge (i.e. +1SD; perceived knowledge score = 7.73 out of 9), ($b = .92$, $SE = .28$),

$z(436) = 3.29$, $p = .001$, 95% CI [.37, 1.46]. Put simply, participants in both conditions chose the more knowledgeable partner but this effect was negligible when participants had low perceived knowledge and was strong when participants had high perceived knowledge (see Figure 4).

Discussion

The Study 3 findings show strong support for the idea that people are uninterested in arguing with someone who has substantially lower knowledge than they do. Perhaps this is because there is little to epistemically gain from a conversation with someone who is very low in issue knowledge—especially given that participants, on average, were above the midpoint on their perceived knowledge ($M = 5.69$ out of 9). Previous research shows that, when people counter-argue against weak arguments or against low-status sources, they do not become any more certain of their beliefs (Tormala & Petty, 2002; 2004). I expected that participant preferences for a relatively more knowledgeable partner would be diminished when the partner had a substantially higher knowledge score (as opposed to one that is similar to the participant). Contrary to my expectations, however, participants preferred a partner with a substantially higher knowledge score than one that had a similar knowledge score to the participant. My primary argument was that people prefer challenging arguments to weak ones only to the degree that they believe they can successfully counter-argue the opposition. This main effect suggests that this preference extends to arguing with others who are ostensibly more knowledgeable on the topic. This is bolstered by previous research that shows that people become more certain after counter-arguing against strong arguments or against high-status sources (Tormala & Petty, 2002; 2004). It is possible that participants in Study 3 had anticipated how certain they might feel after arguing with a more versus less knowledgeable person.

The moderating effects of participant's perceived knowledge provides clear support for my primary argument that people prefer to argue against knowledgeable others but only if they believe they are knowledgeable enough to successfully do so. The main effect of condition discussed above shows that people prefer to argue against knowledgeable others, but the interaction with participant perceived knowledge shows that this effect is non-existent for people low in perceived knowledge and is very strong for people high in perceived knowledge. In terms of Cognitive Energetics Theory, people with high perceived knowledge had a sufficiently high driving force that made them overcome the restraining force of arguing against a relatively more knowledgeable discussion partner. People with low perceived knowledge, however, were constrained by their available knowledge, making them less interested in arguing against a relatively more knowledgeable discussion partner.

Taken together, these findings corroborate and extend research that shows that people who believe they can effectively defend their beliefs are more likely to engage with belief-inconsistent information than people who doubt their ability to defend their beliefs (Albarracín & Mitchell, 2004). The current research builds on this by demonstrating that people seek out more challenging arguments (i.e. by choosing more knowledgeable partners) even if their discussion partner is ostensibly more knowledgeable than the participant. It would be fruitful to explore the mechanism that drives this effect. Previous research would suggest that people who are high in perceived knowledge might engage in an argument with a disagreeing other with expectations that they will leave the conversation feeling even more confident in their beliefs (Albarracín & Mitchell, 2004).

Chapter 5: General Discussion

The research reported in this dissertation provides evidence for the idea that people prefer to counter-argue more difficult messages only if they are high in perceived knowledge. In Study 1, participants who believed they knew little about the carbon tax increased their preference for arguing against the weaker argument when they chose between a moderate and strong argument compared to when they chose between a weak and moderate argument. On the other hand, when participants believed that they knew a lot about the carbon tax, there were no differences in argument preference. Although participants overall preferred to argue against the weaker argument, perceived knowledge of the carbon tax emboldened people to take on stronger arguments. It is important to state, however, that this effect be interpreted with caution as it was the only significant interaction out of four issues.

Study 2 demonstrated support for the idea that perceived knowledge predicts desire to counter-argue more strongly when the argument was strong than when it was weak. This supports the proposition that people who believe they can beat strong arguments will prefer them to weak arguments because they offer a better opportunity to reaffirm one's sense of certainty (Tormala & Petty, 2002; 2004). Here, for the first time, I demonstrate that people can become more interested in attempting to counter *stronger* threats to their beliefs (as opposed to weak threats) so as long as they believe they can defeat them. This is consistent with research showing that people with high confidence that they can defend their attitudes will show smaller information-selection biases in favor of information that supports their beliefs (Albarracín & Mitchell, 2004).

In order to improve upon Study 1, I included in Study 2 a more specific measure that gauges one's perceived ability to counter-argue. The analysis of mechanism demonstrated that people's expectations of effective counter-arguing explain the relationship between argument

strength and choice of whether to counter-argue or ignore the information (i.e. do an unrelated task). These findings support the idea that people prefer direct defenses (e.g. counter-arguing) to indirect defenses (e.g. ignoring, trivialization; Stone, Wiegand, Cooper, & Aronson, 1997), but only if they are sufficiently confident the use of direct defenses will be successful. This is because indirect defenses leave cognitive discrepancies intact whereas direct defenses do not. Thus, the participants in Study 2 likely chose to counter-argue when they believed they could because they were motivated to resolve the inconsistency between their beliefs and the message they read. When they thought they might not be able to counter-argue effectively, they resorted to ignoring the message by choosing to do the unrelated task.

In Study 3, I used a novel paradigm in which people chose among different debate partners who differed in knowledge. The results showed that people largely preferred to debate the more knowledge person. Interestingly, this preference was strengthened when people chose from a relatively more knowledgeable set of potential debate partners—even when one of the debate partners was significantly more knowledgeable than themselves. Supporting my original hypothesis, I found that this effect was non-existent for people with low perceived knowledge and was strong for people with high perceived knowledge. This additionally provides support for the idea that people are willing to engage with difficult tasks if they believe it is the best means to goal attainment (e.g. to bolster one's beliefs; Kruglanski et al., 2012). The results of Study 3 extend this idea by suggesting that arguing against an ostensibly more knowledgeable person is a more ambitious goal for someone who wishes to reaffirm their beliefs. That is, successfully counter-arguing against an expert source makes one more certain of their beliefs than arguing against a non-expert source (e.g. Tormala & Petty, 2002; 2004).

Cognitive Energetics Theory (CET) suggests that people who are high in perceived knowledge have a higher driving force towards defending their beliefs because they assess their ability to be higher than people who are low in perceived knowledge. Thus, they believe they are more equipped to counter-argue belief-threats and should do so as long as their knowledge permits them. People who are low in perceived knowledge presumably believe they are limited in their ability to counter-argue (i.e. low driving force) and therefore their choices are more constrained by the strength of the arguments they encounter (i.e. high restraining force). Indeed, results from Studies 1 and 2 show that people increased their preferences for the easier option when they were given a relatively more difficult set of choices. However, in both studies it was clear that people who believed they knew a lot about the issue showed a smaller preference for taking the easier route than people who believed they knew little about the issue. Put simply, their preferences were less affected by the strength of the belief-inconsistent information because strong arguments do not create as much of a restraining force for someone who is high in perceived knowledge. This fits with the idea that people with limited cognitive resources (i.e. high cognitive load) rely on less informative cognitive shortcuts (e.g. consensus information) but people with available cognitive resources (i.e. low cognitive load) rely on more informative information such as argument content (Pierro et al., 2004). Lowering one's cognitive resources via cognitive load may be functionally equivalent to lowering their perceived knowledge. That is, people do what they must in order to attain their goal (e.g. belief defense) but may be limited by their lack of knowledge or ability (Kruglanski et al., 2012).

Study 3 somewhat diverges from Studies 1 and 2 such that, overall, people preferred the more difficult option (i.e. arguing against a more knowledgeable debate partner). However,

consistent with Studies 1 and 2, people in Study 3 only preferred the more knowledgeable debate partner when they believe they themselves knew a lot about the issue to be discussed.

A plausible reason why the results of Study 3 somewhat diverge from Studies 1 and 2 is the level of uncertainty associated with the belief-threat. When people read the arguments in Studies 1 and 2, they were able to assess whether they could successfully counter-argue them. On the contrary, the paradigm of Study 3 leaves the specific belief-threats a mystery and uses a manipulation of debate partner knowledge. In such a case it is not clear what the other person knows, so participants might assume the partner possesses knowledge of arguments that threaten the participant's position on the issue. CET might suggest that participants assessed the idea of arguing against a knowledgeable partner as a more worthwhile goal. For example, a highly knowledgeable debate partner might offer novel information that challenges the participants' beliefs, which can be later used to further bolster their original position. CET states that people are willing to use more laborious methods of goal attainment, such as arguing against a highly knowledgeable person, if those methods are believed to have the greatest utility among other alternatives. If a person's driving force is sufficiently high, such as having high perceived knowledge on an important issue, then it is plausible to expect that they will subject themselves to strong (and perhaps novel) belief-threats because it ultimately serves their goal of effective belief defense.

Although all three studies provide some support for the hypothesis that people seek to refute strong belief-discrepant information if they are high in perceived knowledge, not all findings were as promising. In Study 1, support for this hypothesis was observed only for the issue of the carbon tax and not for gun control, the junk food tax, or the death penalty. One possibility is that the carbon tax was more important to participants than other issues, which

might explain why there was a heightened desire to counter-argue against information against their beliefs on this issue. However, the data do not support this explanation. Participants rated gun control, the carbon tax, and the death penalty as similarly important ($M_{\text{Gun control}} = 6.52$, $M_{\text{Carbon tax}} = 6.34$, $M_{\text{Death penalty}} = 6.32$) with gun control rated the highest and the junk food tax rated the lowest ($M_{\text{Junk food tax}} = 5.79$). Another possibility is that the issue of the carbon tax was the most novel and unfamiliar issue and thus more susceptible to the experimental manipulation. However, additional data would be required to test this possibility.

The three experiments reported here provide an informative test of Cognitive Energetics Theory as it applies to belief defense. People in these experiments were made to reconcile the driving force towards defending their beliefs and the restraining force of how difficult they would be to defend. These data show that, under specific circumstances (high perceived knowledge in Studies 1 and 3; high expectations of effective counter-arguing in Study 2), people are willing to seek out strong arguments against their beliefs. This is a crucial test of the idea that people might see an epistemically beneficial opportunity in arguing against strong belief-inconsistent arguments, so as long as the arguments do not pose an epistemic risk. If successfully counter-arguing against strong belief-inconsistent information is a more worthwhile endeavor than arguing against weak information (Tormala & Petty, 2002; 2004), then people should seek out such opportunities if they assess their abilities as fit for the task. This is exactly what CET predicts.

A crucial limitation of this research is that there was no assessment of the perceived value (e.g. novelty, potential increase in confidence) of arguing against one message over another or engaging with a particularly knowledgeable (vs. less knowledgeable) debate partner. For example, Albarracín and Mitchell (2004) found that people's expectations that message

information would make them more confident in their beliefs was a primary mechanism driving their selective exposure effects. If, for some people, arguing against *any* belief-discrepant information made them more confident in their beliefs, then their decisions might have been driven by the desire to do the easier task rather than their perceived ability to counter-argue. It would be fruitful for future research to explore the different mechanisms that underlie choice of debate partner versus choice of arguments or tasks.

Still, other knowledge gaps remain. For example, it is unknown whether people's desire to counter-argue message information depends on whether they have encountered such an argument before. It is possible that people desire to counter-argue threats to their beliefs but, after successfully doing so, become less motivated to counter-argue similar arguments. It is also possible that people become more motivated because it is clearer that they will be successful. One way to address this issue is by assessing the novelty of the information to determine its effects on one's desire to counter-argue or ignore it. Although it is plausible for novelty to be a factor in whether people engage with message information (i.e. choose to read or respond to it), other research has failed to find such a relationship (Albarracín & Mitchell, 2004). Nonetheless, it is possible that novelty could have played a role in the current research, especially in Study 3, in which some participants chose to interact with a disagreeing debate partner with substantially higher issue knowledge.

This research calls attention to the importance of focusing on how belief-threatening messages are delivered and who is delivering them. For example, people are quick to doubt an out-group member's reasons for criticizing them or their group, which leads people to substantially reduce their acknowledgement of the criticism (for a review, see Hornsey, 2005). However, if the same exact message comes from an in-group member, people become far more

receptive to the critical information. Thus, it is important to build on the current research by testing if our results change depending one the source of the information.

Conclusion

The way people manage threats to their beliefs, self-interest, or behavior is crucial to adaptive functioning in everyday life. In political discussion, at school, in the workplace, at the doctor, and in close relationships, people will, at some point, be exposed to information that is inconsistent with their beliefs about the world or about themselves. The way people respond to such information is crucial for adaptive functioning because finding ways to ignore or discredit the information prevents oneself from accepting, perhaps badly needed, critical feedback. For example, people respond to negative medical information by questioning the accuracy of the tests as opposed to learning more about the condition with which they had supposedly been diagnosed (Ditto & Lopez, 1992). Additionally, people generally think they are less susceptible than others to widely documented biases such as self-serving biases and the Fundamental Attribution Error (i.e. "bias blind spot"; Pronin, Lin, & Ross, 2002). These defensive responses are key factors in preventing potentially adaptive behaviors (e.g. seeking medical treatment or self-improvement). By downplaying, refuting, or ignoring belief-threatening information, people are preventing themselves from developing more accurate beliefs about the world and themselves.

The research reported here reaffirms the pervasive tendency of people deciding to ignore (rather than confront) criticisms of their political beliefs. However, it has also demonstrated that people are willing to actively counter-argue if they believe they have the knowledge to do so. On one hand, it might not be helpful to our political discourse for people to get increasingly skilled at defending their beliefs. On the other hand, it seems more promising for people to feel

emboldened to engage with belief-threatening information than to ignore it, effectively raising their chances of shifting their views in light of the new information (Albarracín & Mitchell, 2004). If people continue to engage in partisan selective exposure, they will continue to steer clear of potentially valuable criticism. Continuing to fill the knowledge gaps on how people defend their beliefs will be crucial for improving people's receptivity to criticism and fostering constructive public discourse.

Appendix A

Gun Control

Favor Gun Control Strong. Several promising strategies are available to legislators who want to reduce the number of annual gun fatalities in the U.S. The National Institutes of Health (NIH) estimates that a federal law requiring a waiting period of just one day would prevent at least one-third of annual gun suicides, saving more than 7,000 lives every year. Further, studies from the Centers for Disease Control (CDC) show that the majority of homicides are committed with legal firearms (about 68%). The Bureau of Justice Statistics, the CDC, and independent researchers agree that making it more difficult to obtain guns will, even by conservative estimates, save thousands of lives per year.

Favor Gun Control Moderate. Forensic scientists who have studied mass shootings in the U.S. argue that the availability of large magazines to hold large quantities of ammunition increased the number of fatalities over what would have been otherwise possible. The evidence they have uncovered suggests that banning magazines that hold more than 10 rounds would significantly reduce the lethality of typical mass shootings. Implementing a law limiting magazine size would save, on average, about a dozen lives each year, and should be a priority even if doing so is difficult.

Favor Gun Control Weak. The fact that guns are so widely accepted in American culture is harmful and federal law should be used to change this. With the laws as loose as they are now, people are too casual about gun ownership and don't take seriously enough the danger of these weapons. In a discussion with a friend of mine, he told me that he could buy a gun in his own state without even having to go through any training on gun laws or gun usage. We need to limit gun ownership in the U.S. so that we can do away with our gun obsession.

Against Gun Control Strong. A Department of Justice (DOJ) analysis of data from states that have experimented with gun laws suggests that restricting access to guns would probably *increase* the number of gun deaths in the U.S. by about 10%. Scientists that work for the DOJ as well as Centers for Disease Control (CDC) show that 70% of convicted felons reported that they avoided committing crimes when they thought the victim might be armed. The Bureau of Justice Statistics provides evidence that restrictive laws have mostly taken guns away from law-abiding citizens interested in self-defense, but have left the criminal black market unaffected. As a result, ordinary citizens are less able to protect themselves because of the laws and become more vulnerable to criminal violence.

Against Gun Control Moderate. No law would ever be able to take every dangerous weapon out of circulation. In fact, any attempt to restrict gun ownership would be counterproductive. While laws might decrease the total number of guns legally owned by the general public, forensic scientists from the suggest these laws would also increase the black

market trade of guns, which would be harder to regulate and would cause more violence. Gun control laws would also anger large parts of the American public, which is definitely not worth doing for something that probably won't help anyway.

Against Gun Control Weak. Gun control is basically pointless because anyone who uses a gun to kill or injure themselves or another person would almost definitely find another way to do it if they didn't have access to a gun. If you ask a gun control activist to explain concretely how the laws they want would actually help, you'll see that when it comes down to it they don't know what they're talking about and are really just looking for an excuse to limit peoples' constitutional right to bear arms. It's important to put gun control fanatics in their place, even if that means missing a chance to save a few lives.

Death Penalty

Favor Death Penalty Strong. Department of Justice crime data demonstrate that the possibility of a death sentence deters people from committing capital offenses. In states where the death penalty is in practice, there are 20-35% fewer capital offenses per year than in states where the death penalty is not used. While some argue this may result from economic disparities, the fact that the trend is so consistent, regardless of geographical or economic variation, makes it impossible to take this interpretation seriously.

Favor Death Penalty Moderate. The National Center for Criminal Justice conducted systematic interviews with violent offenders that reveal that they often avoid committing crimes that put them at risk of a death sentence, even if they frequently break the law in ways that could earn them lengthy prison sentences. One inmate in Texas who was three years into a 12-year sentence for battery and assault with a deadly weapon explained that he learned his lesson when a member of his gang was convicted of murdering a rival's brother and sentenced to death. While he can't imagine avoiding crime completely when he is released, his gang has already softened their tactics in response to the sentencing and no longer use assassinations to send a message to enemies.

Favor Death Penalty Weak. I asked several of my friends what they thought about the death penalty and they agreed that life in prison without parole is not a harsh enough punishment for the worst criminals, and that the death penalty is the only way to give them what they deserve. Clearly, some crimes are so awful that any punishment less than death would be insufficient. If we, as a society, let murderers off easy because the death penalty makes us uncomfortable, we are nothing more than a nation of cowards.

Against Death Penalty Strong. The Innocence Project at Columbia Law School has documented hundreds of cases of individuals falsely convicted of a crime and sentenced to death. They estimate that as many as 20% of people sentenced to die for a capital crime in the last 50 years were in fact innocent. Even more worrisome, this rate has not declined as forensic technologies have improved over the last two decades. Because of the limitations of DNA and

genetic analyses, the unreliability of eyewitness testimony, and biases that unavoidably affect jury deliberations, states that impose capital punishment are likely to continue the disturbing practice of murdering innocent people that are falsely accused of a crime.

Against Death Penalty Moderate. The most common argument for capital punishment – that it deters violent crime – is not supported by evidence, according to the National Center for Criminal Justice. While advocates of the death penalty cite low capital crime rates in states that practice capital punishment as evidence of its effectiveness, several studies show that not every state that imposes the death penalty has lower rates than average. Furthermore, it is possible that states that enjoy relatively low crime rates owe this to cultural and/or economic, rather than legal, factors. It is therefore unknown whether the possibility of being sentenced to death decreases the incidence of capital crimes. And if we aren't sure that the death penalty will deter future criminals, we shouldn't do it.

Against Death Penalty Weak. It is not acceptable to take another person's life. The reason for this is that killing someone else causes dangerous and irreversible changes in one's mind, so that afterward the killer loses his or her inhibition and is more likely to commit unprovoked violent acts. The same thing happens to American society as long as capital punishment is allowed. By using the death penalty to punish criminals, we become just as bad as them. Therefore, the only way to protect ourselves from being corrupted is to refuse to stoop to the level of violent criminals by punishing them with death.

Junk Food Tax

Favor Junk Food Tax Strong. A study published in the Annual Review of Public Health showed that 97% of cities that have implemented a so-called "junk food tax" saw at least a 26% decline in doctor visits over the first 5 years. Following up on this finding, the American Journal of Medicine published research showing that health care costs have fallen 24% in places that implement the junk food tax compared to a measly 2% over the same period in places that didn't. Meanwhile, healthy packaged foods are becoming increasingly widespread and affordable, offering alternatives to low-income citizens for whom cost is the most important factor driving food choices.

Favor Junk Food Tax Moderate. In recent months several major cities across the country have passed laws implementing a so-called "junk food tax." These measures seem to have broad support, as the Pew Forum published polling data just last week that shows a clear majority of residents in those cities (65%) support the junk food tax. The tax could decrease the number of hospital visits and reduce overall medical costs, and the revenue cities collect from it will enable them to launch a new campaign to promote public awareness of nutrition. The combination of the tax and the campaign should significantly improve health outcomes.

Favor Junk Food Tax Weak. Anyone can tell you that if something costs more money they will buy less of it. If U.S. governments want people to eat healthier, they should implement

a tax on junk food so that people won't buy it anymore. Americans have a reputation around the world for unhealthy eating, and taxing junk foods would pressure restaurants and convenience stores to replace unhealthy options with whole foods like grains and vegetables. This would show the world that Americans care enough about their health to make difficult lifestyle changes when they have to, and that we are just as serious about nutrition as they are.

Against Junk Food Tax Strong. Advocates for a so-called "junk food tax" think it would improve health and reduce medical costs, but a study by the Brookings Institution suggests this policy would actually harm the people it is designed to help. Brookings reports that any tax rate less than 30% would leave the basic economics of the situation as it is, with the least healthy options remaining the cheapest. Given this finding, a tax on junk food would most likely have no effect on the diets of low-income families and individuals, those who are the most vulnerable to dietary-related illnesses, while pushing them even further into poverty. Independent researchers confirm these findings.

Against Junk Food Tax Moderate. A Quinnipiac poll found that a majority of Americans (58%) believe that a junk food tax would fail to reduce health care costs, and an even larger majority (71%) opposes it. While more research by economists and public health experts is still needed, there is reason to believe that the public is right to be skeptical. Tax increases on other products that harm public health, such as cigarettes, have not had much success reducing sales. Unless someone can make a good case for why junk food is different than cigarettes, introducing a special tax would be risky at best.

Against Junk Food Tax Weak. People who want to implement a junk food tax need to mind their own business. Nobody likes taxes unless they are really necessary, but what makes the government an authority on which foods are good for you and which are bad? People's opinions about diet and nutrition are always changing, and 5 years from now people might think the things we are told to avoid now, like sugar, are actually healthy while the things we're told to eat a lot of, like green vegetables, are unhealthy. If we implement a tax on the foods we think are bad now, when the science changes we'll end up taxing the wrong thing.

Carbon Tax

Favor Carbon Tax Strong. The vast majority of scientists agree that carbon emissions contribute to climate change, which they project to displace tens of millions of coastal residents in the coming decades. The U.S. can make a major contribution to slowing this process by following the Department of Energy's (DoE) recommendation to impose a 20% tax on carbon emissions. Experts from the DoE estimate that this would reduce annual emissions by 30-35% in the U.S. and 7-8% globally. It isn't only an environmental issue, the Institute for Economic Policy estimates that the carbon tax will save American taxpayers an average of $18 billion per year by reducing the impact of coastal flooding and natural disasters.

Favor Carbon Tax Moderate. Most scientists from the Academy of Climate Science agree that climate change is happening because of enormous amounts of carbon emissions. The U.S. should take the lead by imposing a 20% tax on carbon emissions. Even if the experts estimates are correct and this step would only reduce global emissions a little (about 2%), other

major countries will likely follow suit. Cumulatively, this would slow the rise of global temperatures and buy time to develop protections against rising sea levels.

Favor Carbon Tax Weak. Americans need to care more about the environment. The U.S. government should sanction anti-environment businesses by imposing a substantial tax on carbon emissions. These businesses are far too profitable and are filled with a bunch of greedy people that don't care about the environment. People who pursue profit without regard to the environment need to finally get a bit of the punishment they deserve. It's not good to consume so much of what this Earth has to provide. We need to get used to consuming less.

Against Carbon Tax Strong. While most scientists agree that our carbon emissions contribute to climate change, economists are pessimistic about the power of any one nation to slow the rate of change without causing intolerable economic harm in the process. The Department of Energy estimates that, to reduce the global emissions rate by just 1%, a carbon tax of more than 50% would be required. The Chamber of Commerce, meanwhile, projects that this would force companies to cut up to 5 million jobs and cause the economy to shrink in ways that are comparable to the Great Depression. This would also include obvious difficulties for the U.S economy such as extremely high gas prices and utility bills.

Against Carbon Tax Moderate. The U.S. should not implement a carbon tax. The Department of Energy estimates that a significant tax on carbon emissions would reduce annual U.S.-based carbon emissions by only very little. It could affect decent-paying jobs and might make gas companies do less business in the U.S. This would not be helpful to the economy. If this tax were implemented, large gas companies might be forced to downsize and significantly cut into their profits.

Against Carbon Tax Weak. Although most scientists agree that our carbon emissions contribute to climate change, imposing a carbon tax would be counterproductive. Taxes are immoral and are just a reason for the government to take money away from people and businesses. The only justifiable ways of dealing with climate change are those that result from market demand, allowing self-interested individuals to compete for the business of people who voluntarily pay for their services. The market does well when government does not get involved.

Appendix B

Death Penalty Knowledge Quiz

How many appeals is a defended given before they are given capital punishment?
- 0
- 1
- 3
- 6
- As many as they want

What is the number of states with the death penalty?
- 15
- 21
- 32
- 41
- 50

What is the number of individuals given the death penalty in the U.S. in 2016?
- 43
- 89
- 115
- 198
- 235

In 2008, states that did not allow capital punishment had a homicide rate that was:
- Significantly lower than the national average (more than 50% lower than the national average)
- Lower than the national average (5-49% lower than the national average)
- The same as the national average
- Higher than the national average (5-49% higher than the national average)
- Significantly higher than the national average (more than 50% higher than the national average)

California's "three-strike law" made it:
- Easier to give the death penalty to repeat offenders – and it was effective.
- Easier to give the death penalty to repeat offenders – and it was ineffective.
- Required life in jail for a three time offender – and is considered effective.
- Required life in jail for a three time offender – and is considered ineffective.
- So that lawyers are not allowed to play baseball.

What is the average cost to taxpayers of life in prison?
- $1 million
- $10 million
- $38 million

- $49 million
- $98 million

The number one comment from victims' families is:
- I felt an overwhelming sense of closure from this event.
- I was surprised by the lack of closure I felt from this event.
- I was not able the guilt I felt because of his/her actions.
- I was relieved and happy that this person got what they deserved.

What is the most accurate definition of capital punishment?
- Physical punishment that involves deliberate infliction of pain as retribution for an offense, or for the purpose of disciplining or reforming a wrongdoer, or to deter attitudes or behavior deemed unacceptable.
- Any punishment in which force is used and intended to cause some degree of pain or discomfort, however light.
- The death penalty is the only example of capital punishment.
- It is a legal term that does not mean anything.
- Any time of punishment that is considered unpleasant by the offender.

What is the most common form of capital punishment given throughout the history of the United States?
- Lethal injection
- Firing squad
- Electric chair
- Hanging
- Poisoning

What percentage of Americans support the death penalty?
- 10%
- 25%
- 40%
- 60%
- 75%

Tables

Table 1a

Mean differences in argument strength ratings

Issue	Condition		t	df
	Weak	Moderate		
Gun control	2.86	4.48	4.18***	41
Junk food tax	3.29	4.21	2.58*	37
Carbon tax	3.27	4.17	3.24**	47
Death penalty	2.60	4.97	6.20***	29

Note. * = $p < .05$, ** = $p < .01$, *** = $p < .001$

Table 1b

Mean differences in argument strength ratings

Issue	Condition		t	df
	Moderate	Strong		
Gun control	3.87	4.84	3.13**	30
Junk food tax	3.86	5.25	3.49***	43
Carbon tax	3.62	4.65	4.25***	35
Death penalty	3.48	4.68	2.37*	25

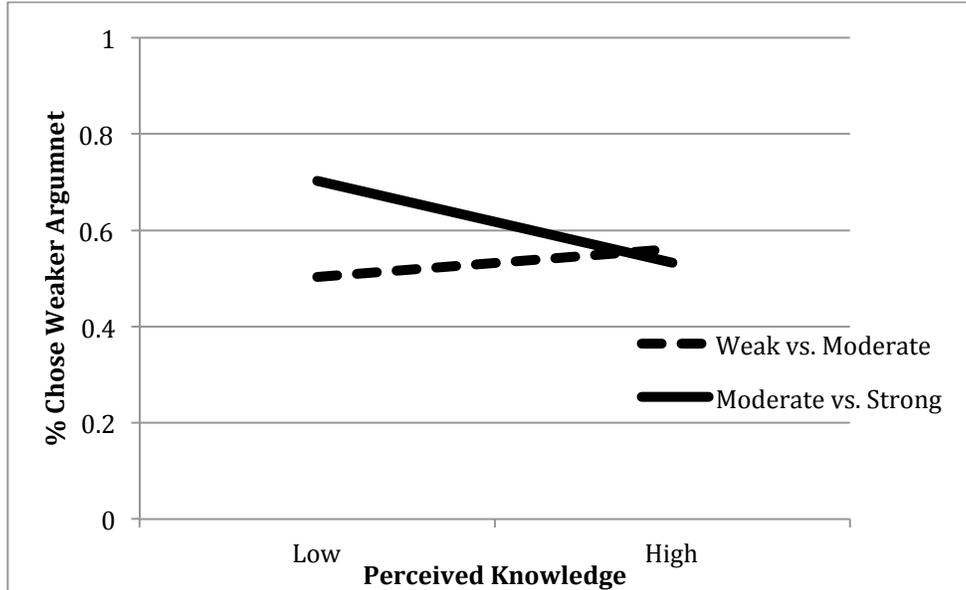*Note.* $* = p < .05$, $** = p < .01$, $*** = p < .001$

Figures

Figure 1. *Percentage of Study 1 participants who chose the relatively weaker argument on the carbon tax as a function of Condition (Weak vs. Moderate; Moderate vs. Strong) and Perceived Knowledge (+/- 1SD).*
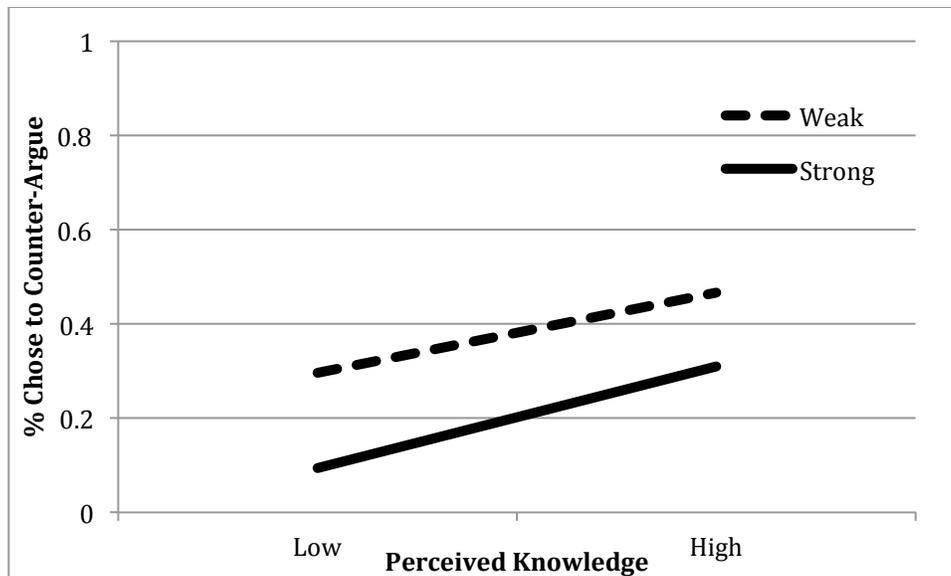
Figure 2. *Percentage of Study 2 participants who chose the counter-arguing task over the unrelated task as a function of Condition (Weak vs. Strong) and Perceived Knowledge (+/- 1SD).*
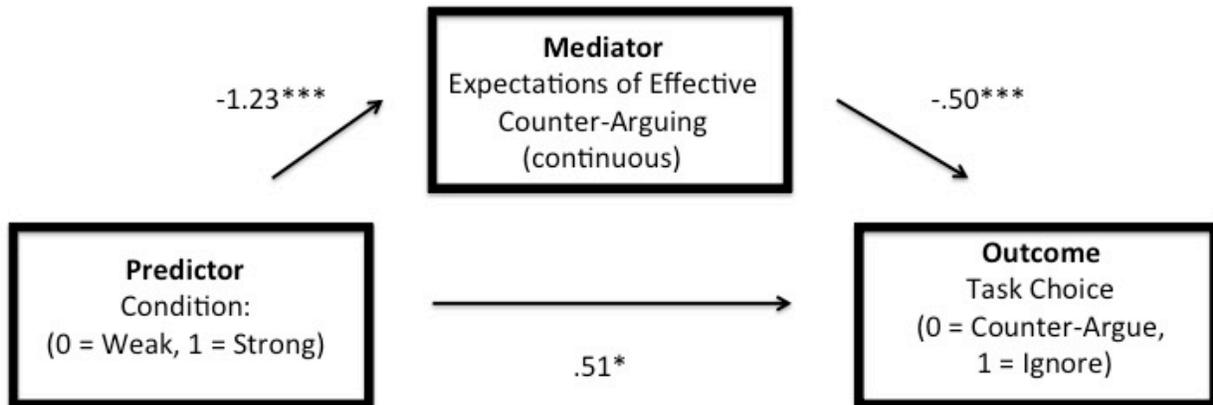
Figure 3. *Mediation model testing the indirect effect of condition (weak vs. strong) on choice of defense strategy (counter-argue vs. ignore) through expectations of effective counter-arguing.* Note. * *p* < .05, *** *p* < .001
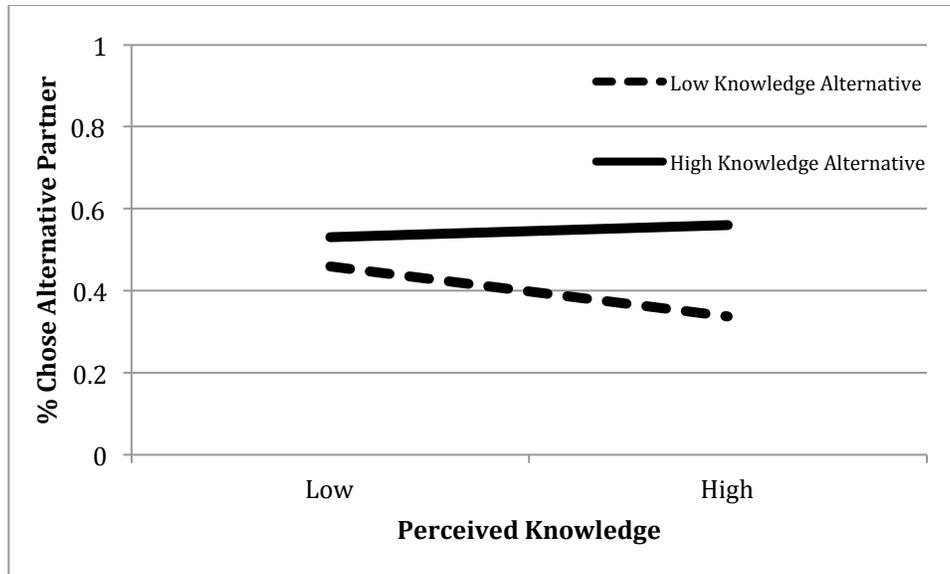
Figure 4. *Percentage of Study 3 participants who chose the alternative partner over one with similar knowledge to themselves as a function of Condition (Low knowledge alternative vs. High knowledge alternative) and Perceived Knowledge (+/- 1SD).*

References

Albarracín, D., & Mitchell, A. L. (2004). The role of defensive confidence in preference for proattitudinal information: How believing that one is strong can sometimes be a defensive weakness. *Personality and Social Psychology Bulletin*, *30*(12), 1565-1584.

Althaus, S. (1998). Information Effects in Collective Preferences. *The American Political Science Review, 92*(3), 545-558.

Arnold, J. R. (2012). The electoral consequences of voter ignorance. *Electoral Studies*, *31*(4), 796–815.

Blanton, H., Strauts, E., & Perez, M. (2012). Partisan Identification as a Predictor of Cortisol Response to Election News. *Political Communication*, *29*(4), 447–460.

Brehm, J. W., & Cohen, A. R. (1962). *Explorations in cognitive dissonance.* San Diego, CA: Academic Press.

Cohen, J. (1988). Statistical power analysis for the behavioral sciences . Hilsdale. *NJ: Lawrence Earlbaum Associates*, *2*.

Ditto, P.H., & Lopez, D.F. (1992). Motivated skepticism: Use of differential decision criteria for preferred and non-preferred conclusions. *Journal of Personality and Social Psychology, 63,* 568–584.

Eagly, A. H., & Chaiken, S. (1993). *The psychology of attitudes.* Fort Worth, TX: Harcourt Brace Jovanovich.

Edwards, K., & Smith, E. E. (1996). A disconfirmation bias in the evaluation of arguments. *Journal of Personality and Social Psychology, 71,* 5-24.

Faul, F., Erdfelder, E., Lang, A. G., & Buchner, A. (2007). G* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior

research methods, *39*(2), 175-191.

Festinger, L. (1957). *A theory of cognitive dissonance.* Stanford, CA: Stanford University Press.

Friesen, J. P., Campbell, T. H., & Kay, A. C. (2015). The psychological advantage of unfalsifiability: The appeal of untestable religious and political ideologies. *Journal of Personality and Social Psychology*, *108*(3), 515–529.

Goldberg, M. H., & Carmichael, C. L. (2017). Language complexity, belief-consistency, and the evaluation of policies. *Comprehensive Results in Social Psychology*, 1-17.

Harmon-Jones, E., & Harmon-Jones, C. (2008). Cognitive dissonance theory: An update with a focus on the action-based model. In J. Y. Shah & W. L. Gardner (Eds.), *Handbook of motivation science* (pp. 71–83). New York: Guilford Press.

Hart, W., Albarracín, D., Eagly, A. H., Brechan, I., Lindberg, M. J., & Merrill, L. (2009). Feeling validated versus being correct: A meta-analysis of selective exposure to information. *Psychological Bulletin*, *135*(4), 555–588.

Hayes, A. F. (2013). *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach*. New York: Guilford Press.

Hornsey, M. J. (2005). Why being right is not enough: Predicting defensiveness in the face of group criticism. *European review of social psychology*, *16*(1), 301-334.

Kiesler, C. A. (1971). *The psychology of commitment.* New York: Academic Press.

Kruglanski, A. W., Bélanger, J. J., Chen, X., Köpetz, C., Pierro, A., & Mannetti, L. (2012). The energetics of motivated cognition: a force-field analysis. *Psychological review*, *119*(1), 1.

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

Lewin, K. (1936). *Principles of topical psychology.* New York, NY: McGraw-Hill.

Litman, L., Robinson, J., & Abberbock, T. (2016). TurkPrime.com: A versatile crowdsourcing

    data acquisition platform for the behavioral sciences. *Behavior Research Methods, 1 - 10.*

Lowin, A. (1967). Approach and avoidance: alternate modes of selective exposure to

    information. *Journal of Personality and Social Psychology*, *6*(1), 1.

Lupia, A. (1994). Shortcuts Versus Encyclopedias: Information and Voting Behavior in

    California Insurance Reform Elections. *The American Political Science Review*, *88*(1),

    63.

Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks:

    Detecting satisficing to increase statistical power. *Journal of Experimental Social*

    Psychology, *45*(4), 867-872.

Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. *Advances*

    in Experimental Social Psychology, 19, 123-205.

Pierro, A., De Grada, E., Mannetti, L., Livi, S., & Kruglanski, A. W. (2004). Bisogno di chiusura

    cognitive e risposta a violazioni normative di carattere quotidiano [Need for closure and

    response to normative violations of everyday nature]. *Giornale Italiano di Psicologia, 31,*

    129 –140.

Pronin, E., Lin, D. Y., & Ross, L. (2002). The bias blind spot: Perceptions of bias in self versus

    others. *Personality and Social Psychology Bulletin*, *28*(3), 369-381.

Rios, K., Goldberg, M. H., & Totton, R. R. (2017). An Informational Influence Perspective on

    (Non) conformity: Perceived Knowledgeability Increases Expression of Minority

    Opinions. *Communication Research*.

Shah, J. Y., & Kruglanski, A. W. (2002). Priming against your will: How accessible alternatives

    affect goal pursuit. *Journal of Experimental Social Psychology*, *38*(4), 368-383.

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

Stone, J., Wiegand, A. W., Cooper, J., & Aronson, E. (1997). When exemplification fails: hypocrisy and the motive for self-integrity. *Journal of Personality and Social Psychology*, *72*(1), 54.

Stroud, N. J. (2008). Media use and political predispositions: Revisiting the concept of selective exposure. *Political Behavior*, *30*(3), 341-366.

Sweeney, P. D., & Gruber, K. L. (1984). Selective exposure: Voter information preferences and the Watergate affair. *Journal of Personality and Social Psychology*, *46*(6), 1208.

Tormala, Z. L., & Petty, R. E. (2002). What doesn't kill me makes me stronger: The effects of resisting persuasion on attitude certainty. *Journal of Personality and Social Psychology*, *83*, 1298-1313.

Tormala, Z. L., & Petty, R. E. (2004). Source credibility and attitude certainty: A metacognitive analysis of resistance to persuasion. *Journal of Consumer Psychology*, *14*(4), 427-442.

Zhou, H., & Fishbach, A. (2016). The pitfall of experimenting on the web: How unattended selective attrition leads to surprising (yet false) research conclusions. *Journal of Personality and Social Psychology*, *111*(4), 493-504.