EFFECTS OF BIASED EXTRAPOLATION ON ATTITUDE EXTREMITY

by

KALEIGH ANN DECKER

Bachelor of Science, 2017
Indiana State University
Terre Haute, Indiana

Master of Science, 2020
Texas Christian University
Fort Worth, Texas

Submitted to the Graduate Faculty of the
College of Science and Engineering
Texas Christian University
in partial fulfillment of the requirements
for the degree of

Doctor of Philosophy
December 2022

EFFECTS OF BIASED EXTRAPOLATION ON ATTITUDE EXTREMITY

By

Kaleigh Ann Decker

Dissertation approved:

C. Lord

[signature]

Cathy Cox

[signature]

L. Brett

L. Brett

For The College of Science and Engineering

ACKNOWLEGEMENTS

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Effects of Biased Extrapolation on Attitude Extremity

Americans are rightly concerned about extreme attitudes in the United States. On issues from immigration to inflation, both sides are adopting more extreme positions. Consider for a moment, topics involving immigration. One side portrays potential immigrants as hard-working families, while the other side portrays them as drug-dealing terrorists. This has sparked interest by researchers and laypeople alike in explanations of how attitudes can become so extreme. The most obvious answer to this question involves an external cause: *accessibility of biased information*. Biased information is more accessible now than in the past because of social and other media networks (Campbell, 2018; Geschke et al., 2019; McCarty, 2019; Parsons & Donehoo, 2019; Sides & Hopkins, 2015). Although I agree that exposure to biased information can certainly make attitudes more extreme, I also believe that focusing only on external causes neglects additional explanations.

A less obvious answer to the attitude extremity mystery involves an internal or psychological cause: *mere thought*. People can think themselves into more extreme positions in the absence of externally provided information (Tesser, 1978). Mere thought can make attitudes more extreme attitudes, for instance, when thoughts are the same valence as the information already known (Tesser & Cowan, 1975; Tesser & Leone, 1977) or when thinkers are confident in their self-generated thoughts (Petty et al., 2002). Relatively little is known, however, about the *process* by which mere thought can make attitudes more extreme. Are there differences between thoughts that go *beyond* what is known and thoughts that simply *review* what is known? Both types of thoughts could be the same valence as the already known information, and the thinker could have confidence in both types of thoughts, but one type might be more likely than the other to make the thinker's attitude more extreme.

Extrapolating, or inferring unknown from known values, is one way of thinking beyond what is known. If someone learns that a group of potential immigrants are *critical* and *fussy,* for instance, they might *extrapolate* beyond the known traits and infer that the group is also *rude*. When asked to evaluate the group, the extrapolated trait might be particularly accessible because of its relevance to the initial information and inform more negative attitudes than attitudes reported by someone who simply *reviewed* the group's known traits. Although this effect of extrapolation on attitude extremity is plausible given previous theory in which attitudes change when associations to the attitude object change (Lord & Lepper, 1999; Schwarz, 2006, 2007; Schwarz & Bohner, 2001; Smith & DeCoster, 2000; Tesser, 1978; Tourangeau, 1992; Wilson & Hodges, 1992), no previously published study has tested the hypothesis that extrapolating beyond what is known is more likely than reviewing what is known to increase accessibility of additional relevant information and to make attitudes more extreme.

Mere Thought and Attitude Extremity

Merely thinking about an attitude object can make attitudes more extreme in the absence of externally provided information. In several studies, Tesser and his colleagues (see Tesser, 1978) gave participants positive or negative information about an attitude object, followed by instructions either to think about the attitude object or to complete an unrelated control task. Participants in Tesser's studies who were instructed to think about the attitude object subsequently reported attitudes more extreme than attitudes reported by control participants. The effect of mere thought on attitude extremity occurred when the attitude object involved a first-person encounter with a confederate who either complimented or criticized the participant (Sadler & Tesser, 1973), a target person or group described only by positive or negative personality traits (Tesser & Cowan, 1975, 1977; Tesser & Leone, 1977), social issues such as

legalizing prostitution and educational grading policies (Tesser & Conlee, 1975), and even inanimate objects like artwork (Tesser, 1976).

The effect of mere thought on attitude extremity tends to be amplified when participants receive ambiguous initial information about the attitude object (Tesser & Cowan, 1977), have a well-developed schema of the attitude object (Tesser & Leone, 1977), or report a greater number of thoughts consistent with the valence of the initial information (Millar & Tesser, 1986; Tesser & Cowan, 1975; Tesser & Leone, 1977). Additionally, when the initial information involves personality traits, the effect of mere thought on attitudes is more pronounced when the initial information involves a smaller than larger sets of traits (Tesser & Cowan, 1975), likely because people use implicit personality theories (Schneider, 1973) to infer additional traits they think an attitude object might also have (Valenti & Tesser, 1981). Tesser (1978) suggested that the opportunity for thought might allow thinkers to reinterpret ambiguous or inconsistent thoughts about the attitude object and to self-generate new thoughts consistent with the valence of their initial attitudes. The process of pruning inconsistent thoughts and generating new, consistent thoughts would result in more extreme attitudes than attitudes reported by participants without a thought opportunity.

Although Tesser's comprehensive research on mere thought laid the foundations for the current research, it did not directly test the present hypotheses. Participants in Tesser's studies were not specifically instructed to *extrapolate* from versus *review* what was known about the attitude object. Instead, participants were simply instructed to *think*. Relatedly, Tesser and his colleagues did not report participants' thoughts during the thought opportunities, so it remains unknown what exactly participants in these studies thought about that made their attitudes more extreme. Participants might have spent time *reviewing* the information they already knew and concluded the information was more extreme *upon further reflection*. Alternatively, participants

might have *extrapolated* from the information they knew by inferring additional attributes they believed the attitude object might also have had. These extrapolated attributes might have been particularly accessible when participants were later asked to evaluate the attitude object, informing more extreme attitudes. Both types of thoughts might be of the same valence as the traits said to describe the attitude object, but thoughts that went beyond what was known might have been more likely than thoughts that focused on what was known to make attitudes more extreme.

Self-Validation & Attitude Extremity

In addition to Tesser's research on the effects of mere thought, extensive empirical evidence from the Elaboration Likelihood Model (ELM, Petty & Cacioppo, 1986) established that the effects of persuasive messages depend importantly on the thoughts that recipients generate for themselves. When people self-generate thoughts consistent with a persuasive message, for instance, their attitudes can become more extreme; however, the opposite occurs when self-generated thoughts are inconsistent with a persuasive message. Self-validation research adds that attitude change also depends importantly on the confidence with which self-generated thoughts are held (Petty et al., 2002). The more confidence people have in their self-generated thoughts about an attitude object, the more validated they believe their thoughts to be, and the more likely they are to use their thoughts to inform subsequent attitudes (Petty et al., 2002; Tormala et al., 2002).

The relationship between thought confidence and attitude change is sensitive to a number of factors. Thoughts about a persuasive message that are easy to generate, for instance, tend to elicit greater confidence and attitude change than thoughts that are difficult to generate (Schwarz et al., 1991; Tormala et al., 2002). In contrast, when people are led to believe that their thoughts might not have been internally generated but instead simply repeat what they heard elsewhere

(Gascó et al., 2018), or are made to doubt the validity of their thoughts on a topic (Clarkson et al., 2013), their self-generated thoughts do not change their attitudes. Ease of thought generation also affects self-generated attitude change through self-validation. People who thought about their personal strengths, for instance, adopted more positive self-attitudes after generating more rather than fewer thoughts, but people who thought about their weaknesses did the opposite (Gandarillas et al., 2018), and people who found it difficult to continue generating thoughts lost confidence in subsequent thoughts (Clarkson et al., 2011; Tormala et al., 2007), which resulted in a lack or reversal of attitude extremity.

The ELM and self-validation research have provided strong evidence that self-generated thoughts, and particularly the confidence with which self-generated thoughts are held, can influence attitudes. The current experiments, however, differ from self-validation studies in notable ways. Like Tesser's research, self-validation research does not typically report participants thoughts during their thought opportunities, nor do they manipulate instructions of *how* participants are supposed to think. Participants in ELM and self-validation studies were typically asked to generate thoughts of *all kinds* rather than using the specific strategy of *extrapolating* from known to unknown values. ELM and self-validation studies tended to measure or manipulate factors related to the source, intended target, or meta-cognitions, rather than measuring the extent to which extrapolating from known information increases accessibility of additional relevant information and makes attitudes more extreme, as suggested by attitude construal theories.

Attitude Construal Theories

The hypothesis that extrapolating from what is known is more likely than reviewing what is known to increase accessibility of additional relevant information and to make attitudes more extreme is supported by construal theories of attitude change. Attitude construal theories posit

that attitudes are informed by the subset of associations that are accessible, or easily obtained, at the time an evaluative judgment is made (Lord & Lepper, 1999; Schwarz, 2006, 2007; Schwarz & Bohner, 2001; Smith & DeCoster, 2000; Tourangeau, 1992; Wilson & Hodges, 1992). This framework suggests that attitude stability and attitude change rely on the same mechanism: accessible associations. Attitudes remain stable, for instance, if similar associations are accessible for the same attitude object across multiple time points or contexts, but attitudes can change when different subsets of associations, or associations of dissimilar valence, are accessible for the same attitude object across multiple time points or contexts (Schwarz & Bless, 2007).

Previous construal theory research examined the types of associations that inform attitudes and the ways in which they do so. Associations that are accessible when evaluating an attitude object can involve category exemplars, personality traits, behavioral information, and the context in which an evaluation is made (Decker & Lord, 2022; Decker et al., in press; Lord & Lepper, 1999; Lord et al., 2004; Schwarz, 2007; Sia et al., 1997). Category exemplars can also prime new associations and attitudes (Lord et al., 2004; Sia et al., 1997), and associations that have been recently or repeatedly activated (Wyer & Srull, 1989) are especially likely to inform attitudes. Empirical research has shown, for instance, that priming participants with the same category exemplar across multiple time points, such as a U.S. president, resulted in attitude consistency, but when a more likeable or less likeable U.S. president was primed, attitudes toward the category "politicians" changed in the corresponding direction (Sia et al., 1997). Finally, if the associations that are accessible when attitudes are first reported match the associations that are accessible during a behavior opportunity, attitude-behavior consistency improves (Ramsey et al., 1994).

The theoretical framework posed by attitude construal theories suggests that *extrapolating* from what is known about an attitude object could make attitudes more extreme in the absence of externally provided information. If attitudes are based on the subset of accessible associations, and recently activated associations are especially likely to inform attitudes, then *extrapolating* from known to unknown attributes will increase accessibility of additional relevant information. These additional relevant associations might subsequently inform new, more extreme attitudes. If, instead, repeated activation of the same associations occurs, such as by *reviewing* attributes known to describe a social group, attitudes will remain relatively stable.

The Present Experiments

In sum, previous research on mere thought, self-validation, and attitude construal theories supports, but has not directly tested the hypothesis that extrapolating beyond what is known about an attitude object is more likely than reviewing what is known to increase accessibility of additional relevant information and to make attitudes more extreme. The present experiments tested the postulated process of attitude change by examining whether *self-generated* extrapolations from the information given activates additional relevant associations and alters attitudes more than *reviewing* the information given (Experiment 1), whether additional methods of extrapolation, such as estimating the *likelihood* of additional traits also inform more extreme associations and attitudes than *reviewing* (Experiment 2), and whether there is a common feature to both methods of extrapolation than can explain the predicted relationship between extrapolation and attitude extremity (Experiment 3).

Attitude object. To test my hypotheses, I used as the attitude object (fictitious) foreign groups who supposedly want to immigrate to the United States. Previous research has used fictitious foreign groups as the attitude object and found that doing so has many benefits (see Decker & Lord, 2022; Decker et al., in press). First, using fictitious foreign groups enhances

experimental control, such that participants begin the experiment without any previous knowledge of the groups. If real groups were used as the attitude object, participants might enter the study with varying degrees of knowledge about the group, which might introduce unnecessary variance in the experimental design. Second, using fictitious foreign groups allowed us to control what participants learned about the group prior to any experimental manipulations. This is especially important so I could ensure that the initial information used to describe the groups was either moderately positive or moderately negative. Third, I believed that using potential immigrant groups as an attitude object increased mundane realism and participant engagement.

Experimental manipulation. As shown in Table 1, the experimental manipulation differed for each Experiment. For Experiment 1, I examined how *self-generated extrapolations* affected associations and attitude extremity. Participants in Experiment 1 learned either moderately positive or moderately negative trait information about two fictitious social groups. For one of the social groups, participants were asked to *self-generate* additional traits they thought the group was likely to have. This way of manipulating extrapolation might be viewed as analogous to giving participants the numbers 2, 4, 6… and asking them to generate additional numbers so as to discover the underlying rule (Wason, 1960). For the other social group, participants were asked to *review* the given traits by typing them in separate textboxes.

Table 1. **Overview of differences among the three experiments.**

	Experiment 1 Self-Generated Extrapolation	Experiment 2 Likelihood Extrapolation	Experiment 3 Types of Traits Extrapolated
Initial Information	Two social groups described by either eight moderately positive or eight moderately negative traits	Two social groups described by either eight moderately positive or eight moderately negative traits	One social group described by either eight moderately positive or eight moderately negative traits
Experimental Manipulation	Extrapolation task: Self-generated eight additional traits for one group Review task: Re-typed initial group information for the other group	Extrapolation task: Rated likelihood of the four most frequently generated traits from Exp. 1 and their antonyms for one group Review task: Re-typed initial group information for the other group	Extrapolation conditions: Rated likelihood of four group traits and their antonyms; Manipulated high vs. low cognitive relevance and high vs. low positivity/negativity of extrapolated traits. No review task.
Potential Mediator	Association extremity: List & rate extremity of the 1st five traits that come to mind when evaluating each group	Association extremity: List & rate extremity of the 1st five traits that come to mind when evaluating each group	Association extremity: List & rate extremity of the 1st five traits that come to mind when evaluating the group
Predicated Attitude Extremity	Extrapolate > Review; Mediation of association extremity	Extrapolate > Review; Mediation of association extremity	Relevant extrapolate > Extreme extrapolate; Mediation of association extremity

Experiment 2 was a direct replication of Experiment 1 except in how participants extrapolated. Instead of self-generating additional traits they thought one of the groups was likely to have, participants in Experiment 2 *estimated the likelihood* of frequently generated additional traits from Experiment 1 participants. This way of manipulating extrapolation might be viewed as analogous to giving participants the numbers 2, 4, 6… and asking them to rate the likelihood that the set of numbers also contains 8.

Experiment 3 assessed whether the extrapolated traits made attitudes more extreme because they were more extreme in their *positivity/negativity* than the traits given or because they were *relevant* to and thus cognitively connected to the traits given. Participants in Experiment 3 learned either positive or negative trait information about one fictitious social group and were asked to extrapolate using likelihood estimates to additional traits that were either high or low in positivity/negativity and either high or low in cognitive relevance to the traits given.

Dependent measures and associations. Following the experimental manipulation, participants in each experiment reported their impressions of the social group(s), support or opposition to the social group(s) immigrating to the U.S., and their willingness to socialize with or do business with group members if they were to immigrate (Decker & Lord, 2022). Experiment 3 included an additional quasi-behavioral measure in which participants were led to believe that their preferences might affect a government agency's immigration decisions. Participants in each experiment also named the first five one-word personality traits that came to mind when they thought of each social group (i.e., associations) and rated the valence of each association they named so that association extremity could be examined as a potential mediator.

Experiment 1: Self-Generated Extrapolation

Experiment 1 assessed whether *self-generated extrapolations* about a social group's positive or negative traits would make subsequent associations and attitudes toward the group more extreme than *reviewing* the initial trait information for another social group. I predicted that participants would report more extreme post-manipulation attitudes and behavioral intentions toward the social group whose initial information they extrapolated than the social group whose initial information they reviewed, regardless of whether the initial group traits were positive or negative. That is, I predicted that those who extrapolated from positive traits would report more positive attitudes and those who extrapolated from negative traits would report more negative attitudes than those who simply reviewed the initial trait information. I further predicted that extremity of associations would mediate the effect of extrapolation on post-manipulation attitudes.

Method

Participants

Using G*Power, an *a priori* power analysis (Faul et al., 2009) based on a 2 (target group: extrapolated group, reviewed group) X 2 (group valence: positive initial traits, negative initial traits) X 2 (order: extrapolation first, review first) X 2 (extrapolated group name: Burum, Z'dura) mixed design analysis of variance (ANOVA), with an estimated effect size of $f = .15$, $\alpha = .05$, power = .80 suggested I would need approximately 360 participants. The initial sample included 430 Amazon Mechanical Turk (MTurk) participants, but 48 participants were excluded from analyses for failing a manipulation check ($n = 21$), not following instructions ($n = 23$; e.g., did not write about the assigned topic), or having duplicate IP addresses ($n = 4$). The final sample consisted of 382 participants (195 female, 185 male, 2 unspecified; age range between 21 and 78,

mdn age = 40, *mdn* years of education = 16). A sensitivity analysis confirmed that the final sample size was large enough to detect an effect at *f* = .14.

Participants qualified for the main experiment if they were 18 years or older, English speaking, and U.S. citizens. Participants also had to indicate that they supported two groups with moderately positive personality traits or opposed two groups with moderately negative personality traits coming to the U.S. This requirement was necessary to ensure participants were responding to the social groups appropriately prior to any experimental manipulations.

Procedure

Qualifying participants first read the following cover letter about the experiment, "We will be giving you information about TWO randomly selected regions from the dozens that have asked the U.S. to recognize them as a country. Although both groups you will read about are being considered for immigration to the U.S., the two groups you will read about were described very differently by unbiased U.S. citizens who lived in each region for over a year. It is very important to make distinctions even between countries that might seem fairly equal in the character of their residents. To help you DRAW DISTINCTIONS between the two groups, we are going to ask you some additional questions about each of the countries, but THE QUESTIONS YOU WILL BE ASKED ABOUT THE TWO COUNTRIES ARE DIFFERENT," (caps in the original cover letter). One of the social groups was said to be from the (fictitious) Burum region of Yemen, and the other group was from the (fictitious) Z'dura region of Algeria.

Initial group description. After reading the cover letter, participants were randomly assigned to read descriptions of either the two moderately positive or the two moderately negative social groups. For the positive social groups, people from the Burum region were described as casual, candid, gregarious, lenient, obliging, congenial, modest, and pardoning, and people from the Z'dura region were described as traditional, methodical, systematic, cultured,

prudent, refined, pragmatic, and straightforward. The traits used to describe the positive groups were rated by 150 MTurk workers in an unrelated study as moderately positive on –5 = *extremely negative for a group of people* to 5 = *extremely positive for a group of people* scales (Burum: *M* likeability = 2.04, *SD* = 0.49; Z'dura: *M* likeability = 2.07, *SD* = 0.54). For the negative social groups, people from Burum were described as fussy, critical, complaining, dissatisfied, finicky, detached, withdrawn, and secretive, and people from Z'dura were described as brash, uncouth, boastful, ostentatious, tricky, showy, overconfident, and possessive. The traits used to describe the negative groups were also rated by 150 MTurk workers as being moderately negative using the same 11-point scales (Burum: *M* likeability = -1.90, *SD* = 0.25; Z'dura: *M* likeability = -1.93, *SD* = 0.37). After reading the social group descriptions, participants completed the experimental manipulation.

Experimental manipulation. Participants were randomly assigned to complete a *self-generated extrapolation* task for one social group and a *review* task for the other group. For the *self-generated extrapolation* task, participants read that my research was interested in personality traits and how easily people could come up with personality traits. They were then instructed to name eight additional personality traits, not mentioned in the initial group description, they thought one of the social groups was likely to have.

For the *review* task, participants were instructed re-type the traits said to describe the other social group in the initial group information. The traits from the initial group information were visible during both tasks. Participants also completed a manipulation check in which they had to indicate whether they were instructed to type new personality traits the social group might have or to re-type the personality traits from the initial social group description. The orders of the self-generated extrapolation and review tasks and the social group whose initial information participants extrapolated or reviewed were counterbalanced.

Associations. After completing the self-generated extrapolation and review tasks, participants named the first five one-word personality traits that came to mind when they thought of each social group (i.e., associations). They also rated the valence of each association they named on –5 = *extremely negative* to 5 = *extremely positive* scales.

Post-manipulation attitudes & behavioral intentions. All participants completed four attitude and behavioral intention items for each social group. Specifically, they reported their impression of the groups from –5 = *extremely negative* to 5 = *extremely positive*, the extent to which they supported or opposed admitting the groups to the United States from -5 = *extremely oppose* to 5 = *extremely support*, their willingness to interact with the groups socially from -5 = *very much avoid* to 5 = *very much approach*, and their willingness to do business with the groups from -5 = *very much avoid* to 5 = *very much approach*. I anticipated that the four items would be highly correlated (Decker & Lord, 2022), and used their average as a measure of participants' overall attitudes in analyses.

Individual difference measures and demographics. Participants completed the Perceived Vulnerability to Disease scale (PVD, Duncan et al., 2009) and the Fear of Covid scale (FOC, Ahorsu et al., 2020) to account for possible effects disease salience from COVID-19 might have had on attitudes. Participants also completed Strahan and Gerbasi's (1972) Social Desirability Scale (SDS) to rule out the possibility that the results were due to experimental demand. Finally, participants reported their age, gender, years of education, marital status, race, ethnicity, income, religious identity, and political ideology[1]. Accounting for these demographic variables has been

---

[1]No demographic variables or scores on the PVD, FOC, or SDS interacted with the effect of extrapolation on attitudes in Experiments 1, 2, or 3, with one exception. In Experiment 2, greater age was related to less extreme attitudes, $b = -0.01$, $SE < .01$, $t = 2.54$, $p = .012$, $R^2 = .01$.

shown to significantly reduce population differences between MTurk and nationally representative samples (Levay et al., 2016).

Debriefing. Finally, participants were debriefed using a debriefing procedure that reliably negates experimentally induced attitude change (Ross et al., 1975).

Experiment 1 Results

Analyses examined whether participants reported more extreme attitudes toward the social group for which they extrapolated additional personality traits than toward the social group for which they reviewed the given traits, and whether extremity of participants' associations mediated the effect of extrapolation on attitude extremity.

Attitudes and behavioral intentions

The conceptual hypothesis was that mere thought is more likely to make attitudes more extreme when individuals with pre-existing attitudes think beyond the information given than when they think about the information given. The conceptual hypothesis would be supported in my specific design if participants reported more extreme post-manipulation attitudes toward the extrapolated than reviewed social group. This hypothesis was tested both separately for each dependent measure and overall, by averaging the four dependent measures.

Separate measures of attitudes and behavioral intentions. Figure 1 shows the results from one-way repeated measures ANOVAs on each of the four dependent measures, with post-manipulation attitudes reversed for the negative groups so that higher numbers represent attitudes that were more extreme (i.e., more discrepant from neutral). As shown in the figure, participants reported more extreme post-manipulation impressions of the extrapolated group ($M_{extremity}$ = 2.70, SD = 0.98) than the reviewed group ($M_{extremity}$ = 2.37, SD = 1.09), $F(1, 381)$ = 41.59, $p <$ .001, d = 0.32. They also reported more extreme post-manipulation attitudes toward admitting members of the extrapolated than the reviewed group to the U.S. ($M_{extremity}$ = 2.82, SD = 1.09 vs.

$M_{\text{extremity}} = 2.57$, $SD = 1.18$), $F(1, 381) = 22.32$, $p < .001$, $d = 0.22$, socializing with them

($M_{\text{extremity}} = 2.52$, $SD = 1.41$ vs. $M_{\text{extremity}} = 2.28$, $SD = 1.49$), $F(1, 381) = 15.61$, $p < .001$, $d = 0.16$, and doing business with them ($M_{\text{extremity}} = 2.49$, $SD = 1.50$ vs. $M_{\text{extremity}} = 2.24$, $SD = 1.54$), $F(1, 381) = 13.31$, $p < .001$, $d = 0.15$. None of these four main effects was qualified by a significant interaction with initial trait valence.

Figure 1. **Attitude and behavioral intention extremity toward the social groups from separate repeated-measures ANOVAs (Experiment 1).**
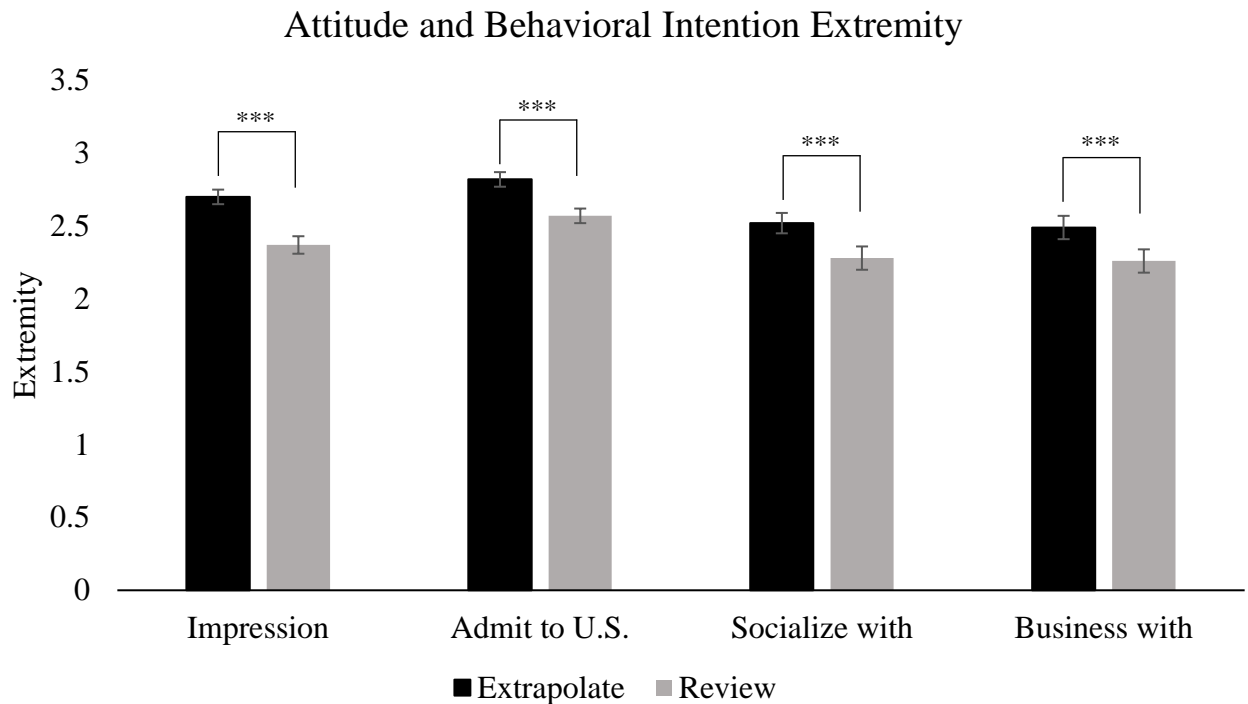
Note: $* = p < .05$, $** = p < .01$, $*** = p < .001$

Overall attitude extremity. The four dependent measures were averaged (negative group $\alpha = .84$; positive group $\alpha = .89$) to form one measure of overall attitude extremity. This overall attitude extremity measure was subjected to a 2 (target group: extrapolated group, reviewed group) X 2 (group valence: positive initial traits, negative initial traits) X 2 (order: extrapolation first, review first) X 2 (extrapolated group name: Burum, Z'dura) mixed model ANOVA, which

yielded a significant effect of target group. Participants reported more extreme post-manipulation overall attitudes toward the extrapolated group ($M_{extremity}$ = 2.63, SD = 1.05) than toward the reviewed group ($M_{extremity}$ = 2.37, SD = 1.15), $F(1, 374)$ = 33.06, $p < .001$, $d = 0.24$. This effect of target group occurred regardless of whether the target group had positive or negative initial traits, interaction $F(1, 374)$ = 1.04, ns, or the extrapolated group's name, interaction $F < 1$. It was qualified, however, by order, interaction $F(1, 374)$ = 16.15, $p < .001$. Greater overall post-manipulation extremity of the extrapolated than reviewed group proved significant for those who extrapolated first (Ms = 2.74 *extrapolation* vs. 2.30 *review*), $F(1, 374)$ = 49.00, $p < .001$, $d = 0.39$, but not for those who extrapolated second (Ms = 2.52 *extrapolation* vs. 2.44 *review*), $F(1, 374)$ = 1.74, ns, $d = 0.08$.

Mediation by associations

Did extremity of participants associations mediate the effect of self-generated extrapolation on attitude extremity?

Using the MEMORE macro for within-subjects mediation (Montoya & Hayes, 2017), I tested whether association extremity mediated the relationship between experimental task (self-generated extrapolation vs. review) and attitude extremity. As shown in Figure 2, participants reported more extreme associations, $b = 0.47$, $SE = .06$, $t = 7.48$, $p < .0001$, and more extreme post-manipulation attitudes toward the extrapolated than the reviewed group, $b = 0.26$, $SE = .05$, $t = 5.72$, $p < .0001$. Association extremity was also related to more extreme post-manipulation attitudes, $b = 0.45$, $SE = .03$, $t = 15.35$, $p < .0001$. Finally, there was a significant indirect effect of self-generated extrapolation on attitude extremity as indicated by a bias-corrected bootstrap (5,000 samples) confidence interval test, $b = 0.21$, 95% CI [.15, .28]. Extrapolating resulted in more extreme associations than reviewing, and the more extreme were associations to the

extrapolated than reviewed group, the more extreme were attitudes toward the extrapolated than reviewed group.

Figure 2. **Effect of self-generated extrapolation on post-manipulation attitude extremity mediated by association extremity (Experiment 1).**
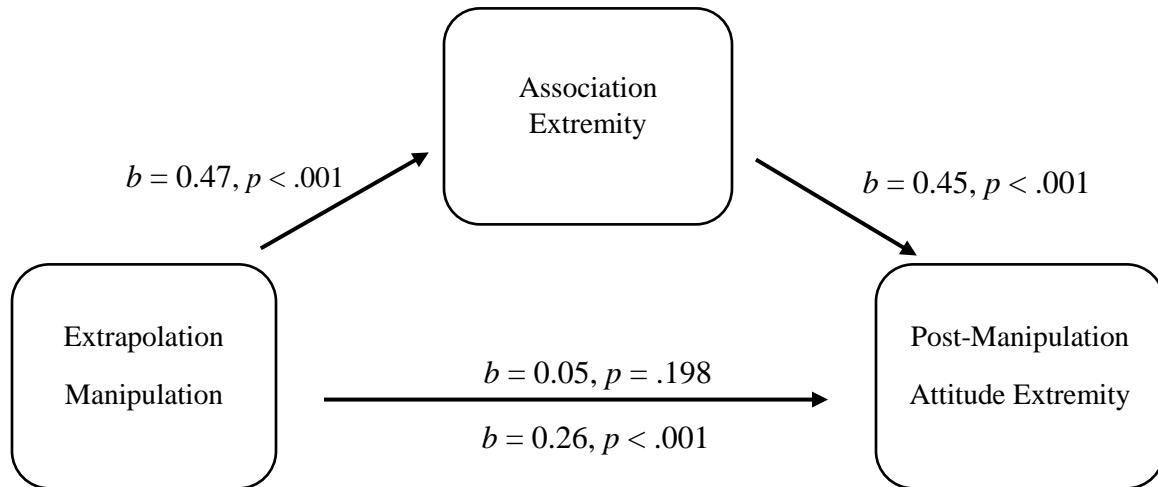
Experiment 1 Discussion

Experiment 1 established that extrapolating beyond what is known about a social group can make attitudes more extreme than reviewing what is known about a social group. Consistent with attitude construal theories (Lord & Lepper, 1999; Schwarz, 2006, 2007; Schwarz & Bohner, 2001; Smith & DeCoster, 2000; Tourangeau, 1992; Wilson & Hodges, 1992), the mediation findings from Experiment 1 also found that the more extreme associations were to the extrapolated than reviewed group, the more extreme were attitudes toward the extrapolated than reviewed group. It remained unclear, however, whether the effects of extrapolation on attitude extremity depended on self-generation. Past research has shown that people are better able to remember self-generated information than information generated by external sources (Slamecka & Graf, 1978), possibly because people tend to confuse the source of information in memory (Johnson et al.,1993), or because people have greater confidence in their self-generated thoughts when they believe their thoughts to be internally generated (Gascó et al., 2018).

Alternatively, there might be *additional* methods of extrapolating that influence associations and attitudes in a comparable way. Self-generating additional traits is not the only way of extrapolating. One might also extrapolate, for example, by estimating the *likelihood* of specific additional traits. After learning that a group is *critical, dissatisfied,* and *detached*, one might estimate a high likelihood that the group is also *rude*, and doing so might increase accessibility of additional relevant associations and ultimately make attitudes more extreme. This seems especially likely because of past research that has shown imagining or explaining the likelihood of even hypothetical events can increase subjective truth value in the event (Anderson et al., 1980; Koehler, 1991). It seemed plausible, then, that extrapolating might influence associations and attitudes not only when self-generating additional group traits, but also when estimating the likelihood of a social group's having additional traits. Experiment 2 assessed this possibility.

Experiment 2: Likelihood Extrapolation

The purpose of Experiment 2 was to determine whether self-generation was necessary or whether additional methods of extrapolation, such as estimating the *likelihood* of a group's having additional unknown traits would make attitudes more extreme than *reviewing* a group's known traits. Experiment 2 employed a procedure identical to Experiment 1, except that the extrapolation task in Experiment 2 involved estimating the likelihood of a social group having the traits most frequently generated by Experiment 1 participants. I predicted that estimating the likelihood of a group's having additional unknown traits would result in more extreme attitudes and behavioral intentions compared to attitudes and behavioral intentions reported after reviewing a group's known traits. As in Experiment 1, the effect of likelihood extrapolation on attitude extremity was predicted to be mediated by more extreme associations to the extrapolated group than reviewed group.

Method

Participants

A total of 429 participants who did not participate in Experiment 1 were recruited from MTurk to participate in Experiment 2. The number of participants was derived from an *a priori* power analysis using G*Power (Faul et al., 2009) based on the same criteria as Experiment 1. Participants had to meet the same qualification criteria as in Experiment 1. A total of 22 participants were excluded from all analyses for not following the instructions during the extrapolation or review tasks. The final sample consisted of 407 participants (214 female, 190 male, 3 unspecified; age range between 18 and 74, *mdn age* = 39, *mdn* years of education = 16). A sensitivity analysis confirmed that the final sample size was large enough to detect an effect at $f = .14$.

Procedure

Qualifying participants followed a nearly identical procedure as in Experiment 1. Participants read either the positive or negative social group descriptions and were randomly assigned to extrapolate for one of the groups and review for the other. The extrapolation task, however, differed from the task used in Experiment 1.

Instead of naming additional traits the social group was likely to have, the extrapolation task for Experiment 2 involved *likelihood estimates*. Participants were asked to rate how likely they imagined it was that the social group had eight additional traits, on –5 = *extremely unlikely* to 5 = *extremely likely* scales. The traits used for the extrapolation task involved the four most frequently generated traits by extrapolators in Experiment 1 and their antonyms (e.g., for the negative social group, they rated the likelihood that the social groups were also *rude* and *polite*, and for the positive social group they rated the likelihood that the groups were also *honest* and *dishonest*). I included antonyms so that both positive and negative traits would be primed by the

manipulation. See Appendix A for extrapolated traits used for each social group with mean extremity of the extrapolated traits and the frequency with which the traits were generated by Experiment 1 participants. The review task was identical to the task used for Experiment 1.

Following the experimental manipulation, all participants listed and rated the valence of five associations to the groups, reported post-manipulation attitudes and behavioral intentions, and completed demographic questions, individual difference measures, and process debriefing (Ross, et al., 1975) using the same measures described in Experiment 1.

Results

As in Experiment 1, analyses examined whether participants reported more extreme attitudes toward the social group for which they extrapolated additional personality traits than toward the social group for which they reviewed the given traits, and whether extremity of participants' associations mediated the effect of extrapolation on attitude extremity. Results are presented first for the separate dependent measures and then for their overall average.

Attitudes and behavioral intentions

Separate measures of attitudes and behavioral intentions. Figure 3 shows the results of one-way repeated measures ANOVAs on each of the four dependent measures, with post-manipulation attitudes reversed for the negative group so that higher numbers represent attitudes that were more extreme (i.e., more discrepant from neutral). As shown in the figure, participants reported more extreme post-manipulation impressions of the extrapolated group ($M_{extremity}$ = 2.68, SD = 0.96) than the reviewed group ($M_{extremity}$ = 2.46, SD = 1.05), $F(1, 406)$ = 16.05, $p <$ .001, d = 0.21. They also reported more extreme post-manipulation attitudes toward admitting members of the extrapolated than the reviewed group to the U.S. ($M_{extremity}$ = 2.82, SD = 1.09 vs. $M_{extremity}$ = 2.63, SD = 1.18), $F(1, 406)$ = 13.71, $p <$ .001, d = 0.16, socializing with them ($M_{extremity}$ = 2.45, SD = 1.50 vs. $M_{extremity}$ = 2.25, SD = 1.61), $F(1, 406)$ = 11.23, $p <$ .001, d =

0.13, and doing business with them ($M_{extremity}$ = 2.48, SD = 1.53 vs. $M_{extremity}$ = 2.34, SD = 1.60), $F(1, 406)$ = 6.24, p = .013, d = 0.09. As in Experiment 1, none of these four main effects was qualified by a significant interaction with initial trait valence.

Figure 3. **Attitude and behavioral intention extremity toward the social groups from separate repeated-measures ANOVAs (Experiment 2).**

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

Overall attitude extremity. The four dependent measures were averaged (negative groups α = .90; positive groups α = .90) to form one measure of overall attitude extremity. This overall attitude extremity measure was subjected to a 2 (target group: extrapolated group, reviewed group) X 2 (group valence: positive initial traits, negative initial traits) X 2 (order: extrapolation first, review first) X 2 (extrapolated group name: Burum, Z'dura) mixed model ANOVA, which yielded a significant effect of target group. Participants reported more extreme post-manipulation overall attitudes toward the extrapolated group ($M_{extremity}$ = 2.61, SD = 1.08) than toward the reviewed group ($M_{extremity}$ = 2.42, SD = 1.16), $F(1, 399)$ = 18.52, p < .001, d = 0.17. This effect

of target group occurred regardless of whether the target group had positive or negative initial traits, interaction $F < 1$, or the extrapolated group's name, interaction $F(1, 399) = 1.17$, *ns*. It was qualified, however, by order, interaction $F(1, 399) = 15.96$, $p < .001$. As in Experiment 1, greater overall post-manipulation extremity of the extrapolated than reviewed group proved significant for those who extrapolated first (Ms = 2.57 *extrapolation* vs. 2.21 *review*), $F(1, 399) = 33.68$, $p < .001$, $d = 0.31$, but not for those who extrapolated second (Ms = 2.64 *extrapolation* vs. 2.62 *review*), $F(1, 399) < 1$, *ns*, $d = 0.02$.

Mediation by associations

Did extremity of participants' associations mediate the effect of likelihood extrapolation on attitude extremity?

As shown in Figure 4, participants reported more extreme associations, $b = 0.21$, $SE = .06$, $t = 3.33$, $p = .0009$, and more extreme post-manipulation attitudes toward the extrapolated than the reviewed group, $b = 0.19$, $SE = .05$, $t = 4.09$, $p = .0001$. Association extremity was also related to more extreme post-manipulation attitudes, $b = 0.44$, $SE = .03$, $t = 15.77$, $p < .0001$. Finally, there was a significant indirect effect of likelihood extrapolation on attitude extremity as indicated by a bias-corrected bootstrap (5,000 samples) confidence interval test, $b = 0.09$, 95% CI [.04, .16]. In line with attitude construal theories and the mediation findings from Experiment 1, extrapolating resulted in more extreme associations than reviewing, and the more extreme were associations to the extrapolated than reviewed group, the more extreme were attitudes toward the extrapolated than reviewed group.

Figure 4. **Effect of likelihood extrapolation on post-manipulation attitude extremity mediated by association extremity (Experiment 2).**
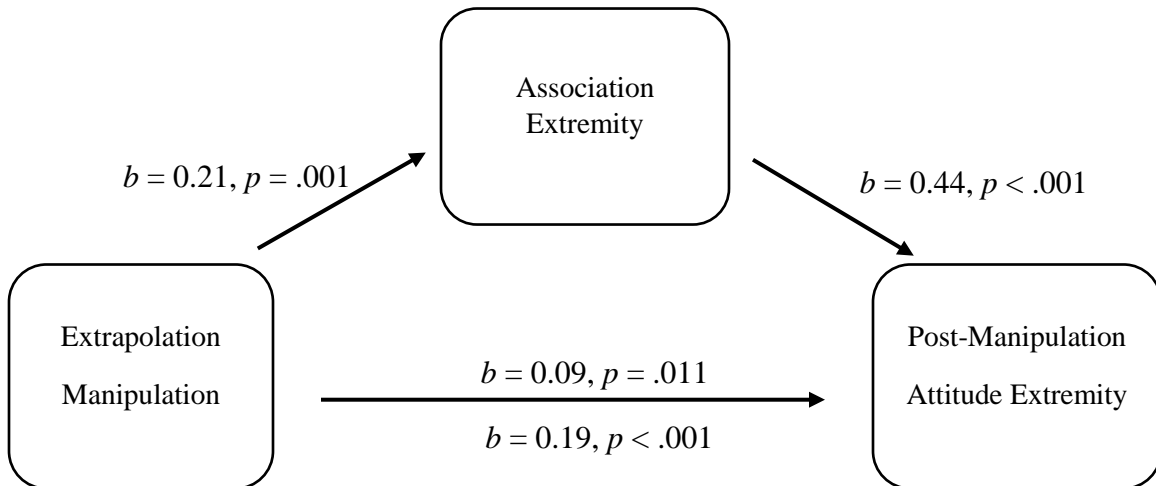
Experiment 2 Discussion

The findings from Experiments 1 and 2 showed that the effect of extrapolation on attitudes occurs when people self-generate additional positive or negative traits they think a social group is likely to have, and when they simply rate the likelihood of a social group having additional negative and positive traits. In both experiments, the effect of extrapolating beyond rather than reviewing the known traits was mediated by more extreme associations to the extrapolated than reviewed social group. It remained unclear, however, whether extrapolation made attitudes more extreme because the traits extrapolated were *more positive/negative* in their valence than the known traits or because the traits extrapolated were *relevant to* and thus cognitively connected to the known traits.

Consider first the idea that extrapolating made attitudes more extreme because the extrapolated traits were *more positive/negative* than the known traits. Participants in Experiment 1 most often named additional traits that were more positive/negative than those given, and participants in Experiment 2 estimated a high likelihood of the social group having those same more positive/negative traits. One possibility, then, is that extrapolation affects associations and

attitudes through *trait positivity/negativity*. As an example of how this might work, most people consider both *rude* and *stupid* as much more negative than *critical* and *complaining*. Extrapolating from *critical* and *complaining* to *rude* and *stupid* constitutes going beyond the information given by focusing on traits farther from neutral, moving attitudes farther from neutral.

Alternatively, extrapolation might affect associations and attitudes through *trait relevance*. The extrapolated traits in Experiments 1 and 2 might have seemed particularly relevant to the social groups' known traits because participants in Experiment 1 generated their own extrapolations and participants in Experiment 2 rated the most frequently generated traits from Experiment 1. Frequently generated extrapolations are by definition relevant to the initial information. As an example of how trait relevance might work, most people consider *rude* to be more relevant than *stupid* to *critical* and *complaining*. Extrapolating from *critical* and *complaining* to *rude* "fits" the theme established by the information given and thus constitutes additional evidence and greater justification for stronger attitudes, in a way that extrapolating to *stupid* does not.

Which is the more important property of extrapolated traits for making attitudes more extreme—how far they are from neutral or how well they "fit" the traits already known? Experiment 3 addressed this research question and tested the hypothesis that relevance (fitting the theme established by the initial information) might prove more important than sheer positivity or negativity.

Experiment 3: Extremity and Relevance of Extrapolated Traits

Experiment 3 assessed whether extrapolating to traits more positive/negative than versus more relevant to the traits known to describe a social group had a greater effect on subsequent associations and attitudes. Participants in Experiment 3 learned positive or negative initial trait

information about one social group and were asked to extrapolate to additional traits that were either high or low in positivity/negativity and either high or low in cognitive relevance to the social group's known traits. I hypothesized that trait relevance would have a greater effect on the relationship between extrapolation and attitude extremity than would trait positivity/negativity. In line with the first two experiments, I further predicted that association extremity would mediate the relationship between relevant trait extrapolation and attitude extremity.

Method

Participants

An *a priori* power analysis using G*Power (Faul et al., 2009) that was based on a 2 (group valence: positive, negative) X 2 (extrapolated trait relevance: high, low) X 2 (extrapolated trait positivity/negativity: high, low) ANOVA with an estimated effect size of $f = .15$, $\alpha = .05$, and power = .80 suggested I would need approximately 351 participants. The initial sample for the 2 X 2 X 2 design included 438 MTurk participants who did not participate in Experiments 1 or 2 (203 female, 233 male, 2 unspecified). An additional 108 MTurk participants (43 female, 63 male, 2 unspecified) who were also not part of any of the previous experiments were subsequently recruited and assigned to a hanging control condition that would be used in secondary analyses[2]. A sensitivity analysis confirmed that the sample size for the main experiment was large enough to detect an effect in a 2 X 2 X 2 ANOVA at $f = .13$.

Participants qualified to participate if they were U.S. citizens and indicated they would have either a moderately negative impression of a moderately negative social group or a moderately positive impression of a moderately positive social group. Impressions were deemed

---

[2] Participant demographics were similar for both samples (main experiment: age range between 19 and 78, *mdn* age = 37.50, *mdn* years of education = 16; hanging control condition: age range between 21 and 69, *mdn* age = 37.00, *mdn* years of education = 16).

moderately negative (positive) if participants reported a -2 or -3 (2 or 3) on a -5 = *extremely negative* to 5 = *extremely positive* scale.

Procedure

Qualifying participants first read either a moderately positive (carefree, congenial, gregarious, lenient, casual, informal, pardoning, and refined; *M* likeability = 1.72, *SD* = 0.68) or moderately negative (brash, possessive, overconfident, uncouth, showy, boastful, tricky, and ostentatious; *M* likeability = -1.93, *SD* = 0.35) description of the Z'dura group of potential immigrants (described in Experiments 1 and 2).

Experimental manipulation. Participants were then randomly assigned to extrapolate by rating the likelihood of one of four sets of four extrapolation traits and their four opposite-valence antonyms. These four trait sets formed a 2 (high vs. low extrapolation trait positivity/negativity) X 2 (high vs. low extrapolation trait relevance) design.

Table 2 shows the extrapolated traits in each cell of the 2 X 2 design, with mean and median positivity and relevance for each cell. High versus low trait positivity/negativity was determined by having 150 MTurk workers in an unrelated study rate group traits on scales from -5 = *extremely negative for a group to be* to 5 = *extremely positive for a group to be.* High versus low relevance was determined by having 200 MTurk workers who were not part of the present experiments rate the likelihood that a social group with the initial group traits would also display each of the extrapolation traits on scales from -5 = *extremely unlikely* to 5 = *extremely likely*.

Table 2. **Mean extremity and relevance of extrapolated traits in the four extrapolation conditions.**

Initial Traits	Relevance	High Extremity	Low Extremity
Positive	High Relevance	Forgiving	Amiable
		Sociable	Mild-tempered
M **Positivity** = 1.72		Happy	Outgoing
		Gracious	Cultured
		M **positivity** = 3.94	*M* **positivity** = 2.72
		Mdn **relevance** = 3.00	*Mdn* **relevance** = 2.50
	Low Relevance	Efficient	Adept
		Self-reliant	Patriotic
		Successful	Clever
		Intelligent	Traditional
		M **positivity** = 3.92	*M* **positivity** = 2.19
		Mdn **relevance** = 1.00	*Mdn* **relevance** = 1.00
Negative	High Relevance	Selfish	Cocky
		Arrogant	Coarse
M **Positivity** = -1.90		Loud	Bragging
		Rude	Stubborn
		M **positivity** = -2.95	*M* **positivity** = -2.23
		Mdn **relevance** = 3.50	*Mdn* **relevance** = 3.50
	Low Relevance	Stupid	Untidy
		Inefficient	Insecure
		Angry	Erratic
		Ignorant	Disorganized
		M **positivity** = -3.37	*M* **positivity** = -2.00
		Mdn **relevance** = 1.00	*Mdn* **relevance** = 1.00

There was also a *hanging control* condition that was run contemporaneously but separately, with different MTurk participants.[3] Participants assigned to the hanging control condition were reminded of either the positive or negative social group's initial traits and immediately reported their attitudes and behavioral intentions toward the group, without the opportunity to extrapolate. The hanging control condition was included to determine baseline attitudes from merely reading the initial information, without further thought.

Attitudes & behavioral intentions. Following the experimental manipulation, all participants completed four attitude and behavioral intention measures. As in Experiments 1 and 2, participants reported the extent to which they supported or opposed admitting the social group to the United States, their willingness to interact with the group socially, and their willingness to do business with the group, on the same scales used in Experiments 1 and 2. The fourth question introduced a new quasi-behavioral dependent measure of interest—how much participants wanted the government to favor immigrants from Z'dura as opposed to immigrants from other geographical regions, from -5 = e*xtremely favor choosing a different group* to 5 = e*xtremely favor choosing Z'dura*. Participants were told I would forward results of this question to the Department of Homeland Security (DHS) so they could use my results as one of the factors they would consider. In other words, participants were led to believe their answers might affect actual immigration decisions.

Associations. After reporting their post-manipulation attitudes and behavioral intentions, participants named the first five associations that came to mind when they thought of the social group and rated the valence of each association. Note that associations were measured before

---

[3] The hanging control conditions were used in secondary analyses only. Participants in the main experiment could not be assigned to the hanging control group, nor could participants in the hanging control group be assigned randomly to any of the four conditions in the main experiment.

attitudes and behavioral intentions in the first two experiments, but associations were measured after attitudes and behavioral intentions in the third experiment.

SDS, demographics, & debriefing. Participants also completed Strahan and Gerbasi's (1972) SDS and reported their age, gender, and years of education[4]. Finally, participants were debriefed (Ross et al., 1975).

Experiment 3 Results

Analyses examined the efficacy of the relevance manipulation, the effects of extrapolating to traits high versus low in cognitive relevance and high versus low in trait positivity/negativity on post-extrapolation attitudes, comparisons to the hanging control condition, and whether extremity of associations mediated the effect of extrapolation on attitude extremity. As in Experiments 1 and 2, results for specific dependent measures and overall attitudes are presented separately.

Relevance manipulation check

The to-be-extrapolated traits used in Experiment 3 were chosen because they had been previously rated by a set of different participants as either relatively high or relatively low in relevance to the traits that would be given for the social group in the main experiment (see Table 2). Did Experiment 3's participants agree? To check the success of my relevance manipulation, I first reversed ratings of antonyms and averaged each participant's likelihood ratings across the extrapolated traits. Participants rated the traits I intended to have high relevance as more likely for the social group ($M_{extremity} = 3.33$, $SD = 0.84$) than the traits I intended to have low relevance ($M_{extremity} = 1.15$, $SD = 1.55$), $t(436) = 18.20$, $p < .001$, $d = 1.75$. It was not only the pre-test

---

[4] The PVD scale (Duncan et al., 2009), FOC scale (Ahorsu et al., 2020), and additional demographic variables from Experiments 1 and 2 were eliminated from Experiment 3 because there was no relationship between post-manipulation attitudes and responses on these items in either of the first two experiments.

participants who viewed the high relevance traits as more likely than the low relevance traits for a group with the given traits. Participants in Experiment 3 agreed.

Attitudes and behavioral intentions

Separate measures of attitudes and behavioral intentions. Figure 5 shows the results from separate 2 (trait positivity/negativity: high vs. low) X 2 (trait relevance: high vs. low) ANOVAs on each of the four dependent measures. As in Experiments 1 and 2, post-manipulation attitudes were reversed for the negative group so that higher numbers represented attitudes that were more extreme. As shown in Figure 5, there was a main effect of trait relevance for three of the four measures: admitting/not admitting to the U.S. (high relevance $M_{extremity} = 2.26$, $SD = 1.20$; low relevance $M_{extremity} = 1.98$, $SD = 1.33$), $F(1, 434) = 5.11$, $p = .024$, $d = 0.22$; doing/not doing business with them (high relevance $M_{extremity} = 2.12$, $SD = 1.67$; low relevance $M_{extremity} = 1.63$, $SD = 1.67$), $F(1, 434) = 9.15$, $p = .003$, $d = 0.29$; and telling DHS to favor/not favor immigrants from Z'dura ($M_{extremity} = 2.42$, $SD = 1.13$ vs. $M_{extremity} = 2.14$, $SD = 1.13$), $F(1, 434) = 6.93$, $p = .009$, $d = 0.25$. In none of the four ANOVAs were trait positivity/negativity or the interaction significant.

Figure 5. **Attitude and behavioral intention extremity toward the social group from separate ANOVAs (Experiment 3).**
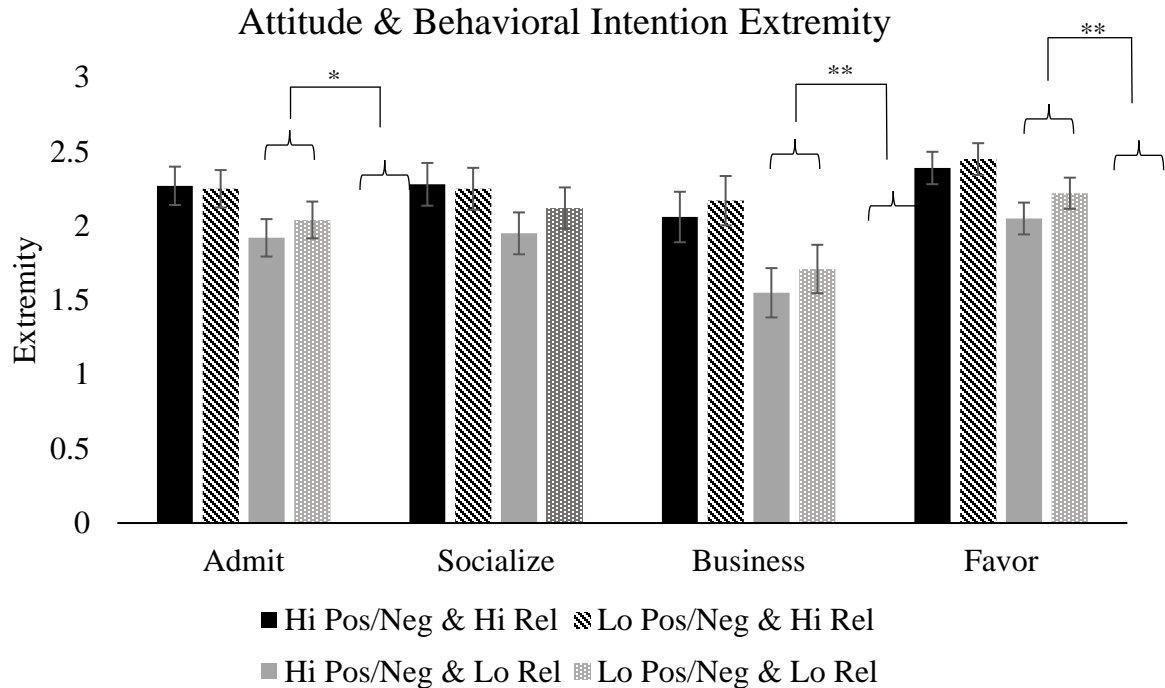
Attitude & Behavioral Intention Extremity

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

Overall attitude extremity. The four dependent measures were averaged (negative group α = .81; positive group α = .76) to form one measure of overall attitude extremity. This overall attitude extremity measure was subjected to a 2 (extrapolated trait relevance: high, low) X 2 (extrapolated trait positivity/negativity: high, low) ANOVA, which yielded only a significant main effect of trait relevance, $F(1, 434) = 9.15$, $p = .003$, $d = 0.28$. Participants who extrapolated highly relevant traits reported more extreme overall attitudes and behavioral intentions ($M_{\text{extremity}}$ = 2.26, SD = 1.08) than did participants who extrapolated to low relevance traits ($M_{\text{extremity}}$ = 1.95, SD = 1.10). The effect of trait positivity/negativity and the two-way interaction were both non-significant, Fs < 1, *ns*.

Comparisons to no-manipulation control. In the main experiment, extrapolating from initial trait information to high relevance traits yielded more extreme post-manipulation attitudes than extrapolating to low relevance traits. Was this because extrapolating to high relevance traits *increased* attitude extremity beyond what would be expected given the extremity of the initial trait information, or because extrapolating to low relevance traits *decreased* attitude extremity from what would be expected given the extremity of the initial trait information? This research question was addressed by planned comparisons to the hanging control condition, in which other MTurk workers contemporaneously reported their attitudes immediately after reading the same initial trait information. Planned comparisons showed that participants in the (averaged) two high relevance extrapolation groups reported more extreme overall attitudes ($M_{extremity}$ = 2.26, SD = 1.04) than did participants in the hanging control condition ($M_{extremity}$ = 1.93, SD = 1.21), $t(541)$ = 2.53, p = .012, d = 0.33. Participants in the (averaged) two low relevance extrapolation groups did not ($M_{extremity}$ = 1.95, SD = 1.10, $t < 1$). These comparisons, although not definitive, suggested that extrapolating to low relevance traits left attitudes unchanged, whereas extrapolating to high relevance traits made attitudes more extreme.

Mediation by associations

Did extremity of associations mediate the effect of high vs. low relevance extrapolation on attitude extremity?

Mediation, using the PROCESS macro for between-subjects mediation (Hayes, 2022), was assessed by comparing the high vs. low relevance extrapolation conditions as the IV, association extremity as the mediator, and overall attitude extremity as the DV. As shown in Figure 6, *high relevance* extrapolation conditions were associated with more extreme associations, b = 0.32, SE = .11, t = 2.96, p = .003, and more extreme attitudes than *low relevance* extrapolation conditions, b = 0.20, SE = .09, t = 2.20, p = .028. Association extremity

was also related to more extreme attitudes, $b = 0.36$, $SE = .04$, $t = 9.03$, $p < .0001$. Finally, there was a significant indirect effect of *high relevance* extrapolation on attitude extremity as indicated by a bias-corrected bootstrap (5,000 samples) confidence interval test, $b = 0.12$, 95% CI [.04, .20]. Extrapolating traits *high in cognitive relevance* resulted in more extreme associations to the social group than extrapolating traits *low in cognitive relevance*, and more extreme associations was related to more extreme attitudes.

Figure 6. **Effect of high vs. low relevance extrapolation on post-manipulation attitude extremity mediated by association extremity.**

Association Extremity

$b = 0.32$, $p = .003$

$b = 0.36$, $p < .001$

Hi vs. Lo Relevance Extrapolation Conditions

$b = 0.32$, $p = .001$

$b = 0.20$, $p = .028$

Post-Manipulation Attitude Extremity
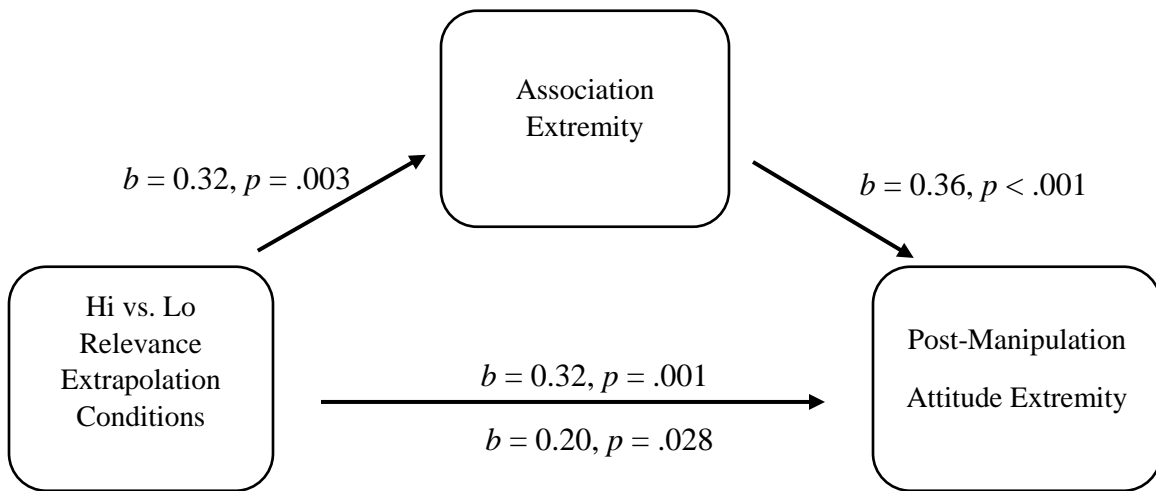
Experiment 3 Discussion

The findings from Experiment 3 showed that extrapolating to highly relevant traits resulted in more extreme attitudes and behavioral intentions than extrapolating to low relevance traits. The effect of high relevance trait extrapolation on attitudes occurred regardless of whether the traits participants extrapolated were high versus low in trait positivity/negativity. As in the previous two experiments, extremity of the associations that came to mind following the extrapolation task mediated the effect of extrapolation on attitudes. Specifically, extrapolating traits high in relevance to the social group's initial information resulted in more extreme

associations to the social group than extrapolating traits low in relevance, and more extreme associations was related to more extreme attitudes.

General Discussion

The findings from the three experiments supported the hypothesis that extrapolating beyond what is known is more likely than reviewing what is known to increase accessibility of additional relevant information and make attitudes more extreme. This pattern of results occurred when initial attitudes were moderately positive or moderately negative, regardless of the social group participants extrapolated or reviewed. It also occurred for different methods of extrapolation, such as when participants self-generated their own extrapolations (Experiment 1) or rated the likelihood of traits frequently extrapolated by others (Experiment 2). Experiment 3 showed also that the effect of extrapolation on attitudes was greatest when participants extrapolated traits relevant to the traits said to describe a social group, regardless of extrapolated trait extremity. Finally, each experiment yielded consistent evidence that the effect of extrapolation on attitude extremity was mediated by more extreme accessible associations to the extrapolated than reviewed social group (Experiments 1 and 2) or when the extrapolated traits were highly relevant to the group's known traits (Experiment 3). The findings from these experiments are consistent with and extend past research and theory on several topics.

The current experiments replicated and extended Tesser and colleagues' (Millar & Tesser, 1986; Sadler & Tesser, 1973; Tesser & Conlee, 1975; Tesser & Cowan, 1975, 1977; Tesser & Leone, 1977; Valenti & Tesser, 1981) research on the effects of mere thought on attitude extremity. Their research showed that simply thinking about an attitude object could make the thinker's attitude more extreme in the absence of externally provided information. The current experiments extended this research by finding that *certain types* of thought are more likely than others to make attitudes more extreme. I instructed participants in Experiments 1 and

2 to think about two social groups by *extrapolating* from known to unknown personality traits for one group, and by *reviewing* the known personality traits for the other. Post-manipulation attitudes were more extreme toward the social group for which participants extrapolated traits than toward the group for which they reviewed the initial traits. This effect of thought type occurred when participants in Experiment 1 self-generated additional traits and when participants in Experiment 2 estimated the likelihood of a group having additional traits.

The three experiments also provided evidence for the generality of the effects of extrapolation on extreme attitudes by establishing that the effects work similarly for moderately positive and moderately negative initial attitudes. This finding is consistent not only with Tesser's (1978) research on mere thought, but also with research on the ELM. The ELM holds that self-generated thoughts can make attitudes more extreme in the absence of externally provided information, regardless of the attitude object's valence (Petty & Cacioppo, 1986). In all three of the current experiments, extrapolating made initial moderately positive and initial moderately negative attitudes more extreme, an effect not qualified by demographic variables, by social desirability (Strahan and Gerbasi, 1972), or by disease salience from Covid-19 (Ahorsu et al., 2020; Duncan et al., 2009). This generality suggests that the process involved in extrapolation might apply more generally to other attitude objects or other attributes, and perhaps especially for people with moderate initial attitudes.

The present findings are also consistent with and extend construal theories of attitude change that suggest attitudes are informed by the associations that are accessible during evaluation (Lord & Lepper, 1999; Schwarz, 2006, 2007; Schwarz & Bohner, 2001; Smith & DeCoster, 2000; Tourangeau, 1992; Wilson & Hodges, 1992). In all three experiments, extremity of the associations that came to mind following the extrapolation manipulations mediated the relationship between extrapolating vs. reviewing (Experiments 1 and 2) and high vs. low

relevance extrapolation (Experiment 3) and attitude extremity. These findings extend construal theories in two notable ways. First, I found that association extremity mediated the effect of extrapolation on attitudes regardless of whether associations were measured before (Experiments 1 and 2) or after (Experiment 3) participants reported their attitudes and behavioral intentions, which suggests that the order in which associations and attitudes are measured might be interchangeable. Second, the relevance findings from Experiment 3 suggest that attitudes are informed not only by accessible associations, but also by associations that seem particularly *relevant* to the attitude object and the evaluation.

The importance of trait relevance in Experiment 3 also fits well with research on implicit personality theories. Implicit personality theories suggest that people draw inferences from and inflate the relationship between personality traits (Schneider, 1973). Experiment 3 found that the effect of extrapolation on associations and attitudes was more likely to occur when participants extrapolated traits high rather than low in cognitive relevance to the group's known traits. Participants might have found it easier to draw connections among traits that seemed highly relevant to the traits they knew compared to traits that seemed less relevant, especially because people often expect different personality traits to coincide (Bruner et al., 1958; Lay & Jackson, 1969). Relatedly, participants in Experiment 3 might have considered the extent that the extrapolated traits seemed relevant to the *gist* of the initial traits and believed the high-relevance traits to be conceptually related (Reyna, 2012). People often rely on gist information when they are asked to draw inferences (Reyna & Brainerd, 1995).

Extrapolating is likely only one of many types of thought that enables more extreme associations and attitudes. Other research has found, for instance, that generalizing about a group's attributes across empirically distinct settings can also make attitudes more extreme in the absence of new information (Decker & Lord, 2022). In this research, participants were asked to

generalize about the likelihood that a social group would express their assigned traits in settings that were conceptually distinct from those referenced in the initial information (Kenrick et al., 1990). Consistent with a correspondence bias (Jones, 1990), participants thought it highly likely that the generalized group would express the assigned traits in additional settings, and they subsequently reported more extreme attitudes toward the group. These studies differed from the current experiments, however, because the manipulation in the generalization studies involved drawing inferences from a group's traits in *one type of setting* to the same traits in a *different type of setting*, whereas the manipulation in the current experiments involved drawing inferences from a group's traits *in general* to the group's additional traits *in general*.

Other research has found that writing persuasive social media posts to friends and unbiased strangers can also make attitudes more extreme in the absence of new information (Decker et al., in press). This research had non-neutral participants write social media posts about a social group said to have either moderately positive or moderately negative traits. In writing their social media posts, participants inserted statements that went beyond what they knew about the social group and subsequently reported more extreme attitudes than did participants who wrote on another topic. Unlike the communication experiments, the current experiments involved specific instructions to either self-generate additional traits or estimate the likelihood of a social group having additional traits, without instructions to be persuasive. The current experiments also aimed to study extrapolation *in general* and was not confined to an online communication context. Participants in the current experiments were given no reason to believe that their extrapolations would be publicly available or that they might change anyone else's opinion.

Limitations

Generality of the present findings was limited in several ways. One limitation involved the order effects found in Experiments 1 and 2 in which attitudes were more likely to be extreme toward the extrapolated than reviewed group when participants extrapolated first than when they reviewed first. Order was counterbalanced to control for possible fatigue effects in which participants might exert more mental effort on whichever task came first. It might be tempting to explain the task X order interaction in Experiments 1 and 2 by reference to fatigue (Webster, Richter, & Kruglanski, 1996), except that extrapolation means were just as extreme when participants extrapolated second as when they extrapolated first. Review condition means, in contrast, were less extreme when participants reviewed second than when they reviewed first. Fatigue might be blamed for the pattern of review condition means, but not for the pattern of extrapolation condition means, nor for the significant interaction.

A more convincing interpretation might be that participants reported equally extreme attitudes for the group considered first because, in such novel tasks, they needed first to establish an adaptation level anchor (Helson, 1948), which was subject to ceiling effects (Wang et al., 2008). Having reported extreme attitudes after extrapolating for the first group, it might have become evident to participants that the "evidence" for the reviewed group did not seem as strong. When they completed the two tasks in the opposite order, participants set their anchor so extreme for the reviewed group that they could not justify reporting greater extremity for the extrapolated group. This adaptation level problem might be averted in future research by having participants complete both the extrapolation and the review tasks first, before reporting their attitudes toward both groups—a procedure that was not used in the present experiments for fear participants would confuse the two groups (with unfamiliar names) when reporting their attitudes. Finally, it should be noted that means for all four experimental (extrapolation) conditions in Experiment 3

were more extreme than in the hanging control or "no thought" condition. Because Experiment 3 lacked a review condition, it remains possible that reviewing the initial information makes attitudes more extreme than merited by that information but extrapolating beyond the initial information generates even greater attitude extremity.

The present results also afforded only indirect evidence of how thought confidence might contribute to the effects of extrapolation on associations and attitudes. Self-validation research has found that people believe their self-generated thoughts to be especially credible, and the greater confidence they have in their thoughts, the more likely their thoughts are to inform their attitudes (Petty et al., 2002; Tormala et al., 2002). In the current research, self-generated extrapolations and estimates of trait likelihood made positive and negative attitudes more extreme than reviewing, and although thought confidence was not measured in the reported experiments, participants' likelihood estimates from the extrapolation tasks in Experiments 2 and 3 might be thought of as a reasonable substitute for confidence ratings. People tend to believe even hypothetical information they explain or imagine (Anderson, 1983; Anderson et al., 1980; Koehler, 1991), so perhaps extrapolating increased the subjective truth value of the extrapolated traits, and thus, their confidence in the likelihood that the group would display the extrapolated traits.

A third limitation involved the mediation results from the three experiments. Consistent with attitude construal theories (Lord & Lepper, 1999; Schwarz, 2006, 2007; Schwarz & Bohner, 2001; Smith & DeCoster, 2000; Tourangeau, 1992; Wilson & Hodges, 1992), association extremity mediated the effect of extrapolating on attitude extremity, such that participants reported more extreme associations after extrapolating than after reviewing (Experiments 1 and 2) or after extrapolating to traits high vs. low in cognitive relevance to the given traits, and these more extreme associations were related to more extreme attitudes. Importantly, mediation

findings are purely correlational and do not imply causation (Spencer et al., 2005). Other

research has manipulated accessible associations via priming, however, and found that priming

different exemplars across multiple time points can change participants' associations and

attitudes (Lord et al., 2004; Sia et al., 1997). Although I did not manipulate participants'

associations in the current research, I did find that associations mediated the effect of

extrapolation on attitudes regardless of whether associations were measured before or after

attitudes and behavioral intentions.

Future Directions

The current experiments examined different methods of extrapolation, common features of

self-generated and likelihood extrapolations, and possible mechanisms underlying these effects.

Although the current experiments advanced knowledge of the processes underlying self-

generated attitude change through extrapolation, there are many promising future directions

worth investigation. The attitude object in all three experiments, for example, involved fictitious

foreign groups described by moderately positive or moderately negative personality traits. It

might be that extrapolation works best with this type of initial information and attitude object but

not with others, possibly because people use implicit personality theories to draw inferences

from personality traits (Schneider, 1973), and drawing inferences from information about foreign

groups might seem especially relevant or easy (Schwarz et al., 1991) because of increased

national attention to immigration topics.

Future research should explore whether the effects of extrapolation on associations and

attitudes would generalize, for instance, if a different group or organization was described by a

different type of initial information, such as exemplars, goals, or intentions. Some people

intensely like or dislike organizations like Planned Parenthood or the National Rifle Association,

not necessarily because of known or inferred traits, but because of their *goals, intentions,* and

expressed values. Relatedly, extrapolating beyond what is known about an *individual*, such as in instances of online dating in which an individual infers additional attributes based on a potential mate's dating profile, might also make associations and attitudes more extreme. Extrapolating in this context might even increase or decrease willingness to message or go on a date with a potential mate. Future research should also explore whether extrapolating from attributes said to describe an *object*, such as in instances of product development, would change consumers attitudes and willingness to purchase or use a product. One might argue that extrapolating from known to unknown attributes of *any type* could make attitudes more extreme, but this conclusion cannot be drawn until additional research is conducted.

In addition to exploring the effect of extrapolation on attitudes toward other types of attitude objects described by other types of initial information, future research should also try to identify possible individual differences that might moderate these effects. In the current experiments, I examined whether the effect of extrapolation on attitudes was qualified by a range of demographic variables or individual difference measures related to disease salience from COVID-19 (Ahorsu et al., 2022; Duncan et al., 2009). I did not, however, measure possible relevant individual differences in need for cognition (Cacioppo & Petty, 1982), need for affect (Maio & Esses, 2001), open minded thinking (Svedholm-Häkkinen & Lindeman, 2018), or stable versus malleable attitudes (Akhtar & Wheeler, 2016; Petrocelli et al., 2010). Moderation of individual differences in need for cognition, for instance, might point to self-generation of likely additional attributes as instrumental. Relatedly, there might be individual differences that moderate the effect of extrapolation on attitudes only for specific types of attitude objects or initial attitudes of a specific valence, instead of moderating the process *in general*. Individual differences in judgments of moral character (Bocian et al., 2018) or belief that the world is a

dangerous place (Altemyer, 1988), for example, might moderate the effect of extrapolation on negative attitudes toward outgroups but not positive attitudes toward ingroups.

Concluding Remarks

The current research was inspired by growing concern about extreme attitudes in the United States. Many researchers and laypeople have blamed the rise of extreme attitudes on biased information found in social and other media, such as inflammatory social media posts (McCarty, 2019; Parsons & Donehoo, 2019), "filter bubbles" (Geschke, et al., 2019), and exposure to attitude inconsistent information online (Bail et al., 2018). Focusing only on the ways in which encounters with *new* information influences attitudes ignores an equally important process contributing to this global challenge: *self-generated attitude change*. The current findings suggest that the mere act of extrapolating from known to unknown values, without exposure to any new or potentially biased information, can also make attitudes more extreme. Extrapolating is but one type of thought that can make attitudes more extreme, and there are likely many others that yield similar effects on attitudes. I cannot hope to minimize attitude extremity until researchers prioritize attitude extremity from both exposure to new (possibly biased) information, as well as attitude extremity created from our own thoughts.

References

Ahorsu, D. K., Lin, C. Y., Imani, V., Saffari, M., Griffiths, M. D., & Pakpour, A. H. (2020).

The fear of COVID-19 scale: Development and initial validation. *International Journal of Mental Health and Addiction, 20*(3), 1537–1545. https://doi.org/10.1007/s11469-020-00270-8

Akhtar, O., & Wheeler, S. C. (2016). Belief in the immutability of attitudes both increases and decreases advocacy. *Journal of Personality and Social Psychology, 111*(4), 475–492. https://doi.org/10.1037/pspa0000060

Altemeyer, B. (1988). *Enemies of freedom.* San Francisco, CA: Jossey-Bass.

Anderson, C. A., Lepper, M. R., & Ross, L. (1980). Perseverance of social theories: The role of explanation in the persistence of discredited information. *Journal of Personality and Social Psychology, 39*(6), 1037–1049. https://doi.org/10.1037/h0077720

Anderson, C. A. (1983). Imagination and expectation: The effect of imagining behavioral scripts on personal influences. *Journal of Personality and Social Psychology, 45*(2), 293–305. https://doi.org/10.1037/0022-3514.45.2.293

Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M., Lee, J., Mann, M., Merhout, F., & Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences of the United States of America, 115*(37), 9216–9221. https://doi.org/10.1073/pnas.1804840115

Bocian, K., Baryla, W., Kulesza, W. M., Schnall, S., & Wojciszke, B. (2018). The mere liking effect: Attitudinal influences on attributions of moral character. *Journal of Experimental Social Psychology, 79*, 9–20. https://doi-org.ezproxy.tcu.edu/10.1016/j.jesp.2018.06.007

Bruner, J. S., Shapiro, D., & Tagiuri, R. (1958). The meaning of traits in isolation and in combination. In R. Tagiuri and L. Petrullo (Eds.), *Person perception and interpersonal behavior*, (pp. 276-288). Stanford: Stanford University Press.

Cacioppo, J. T., & Petty, R. E. (1982). The need for cognition. *Journal of Personality and Social Psychology, 42*(1), 116–131. https://doi.org/10.1037/0022-3514.42.1.116

Campbell, J. E. (2018). *Making sense of a divided America*. Princeton University Press.

Clarkson, J. J., Tormala, Z. L., & Leone, C. (2011). A self-validation perspective on the mere thought effect. *Journal of Experimental Social Psychology, 47*(2), 449-454. doi:10.1016/j.jesp.2010.12.003

Clarkson, J. J., Valente, M. J., Leone, C., & Tormala, Z. L. (2013). Motivated reflection on attitude-inconsistent information: An exploration of the role of fear of invalidity in self-persuasion. *Personality and Social Psychology Bulletin, 39*(12), 1559–1570. https://doi.org/10.1177/0146167213497983

Decker, K. A. & Lord, C. G. (2022). Generalizing across behavior settings can make attitudes toward social groups more extreme in the absence of new information. *European Journal of Social Psychology, 52*, 91– 104. https://doi.org/10.1002/ejsp.2817

Decker, K. A. & Lord, C. G., & Holland, C. (in press). Communicating beyond the information given can make the communicator's attitudes toward a social group more extreme. *Journal of Social Psychology*. https://doi.org/10.1080/00224545.2022.2133678

Duncan, L. A., Schaller, M., & Park, J. H. (2009). Perceived vulnerability to disease: Development and validation of a 15-item self-report instrument. *Personality and Individual Differences, 47*(6), 541–546. https://doi.org/10.1016/j.paid.2009.05.001

Faul, F., Erdfelder, E., Buchner, A., & Lang, A. (2009). Statistical power analyses using

    G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research*

    Methods 41, 1149–1160. https://doi.org/10.3758/BRM.41.4.1149

Gandarillas, B., Briñol, P., Petty, R. E., & Díaz, D. (2018). Attitude change as a function of the

    number of words in which thoughts are expressed. *Journal of Experimental Social*

    Psychology, 74, 196-211, https://doi.org/10.1016/j.jesp.2017.09.012

Gascó, M., Briñol, P., Santos, D., Petty, R. E., & Horcajo, J. (2018). Where did this thought

    come from? A self-validation analysis of the perceived origin of thoughts. *Personality*

    and Social Psychology Bulletin, 44(11), 1615–1628.

    https://doi.org/10.1177/0146167218775696

Geschke, D., Lorenz, J., & Holtz, P. (2019). The triple-filter bubble: Using agent-based modeling

    to test a meta-theoretical framework for the emergence of filter bubbles and echo

    chambers. *British Journal of Social Psychology, 58*(1), 129-149.

    https://doi.org/10.23668/psycharchives.2541

Helson, H. (1948). Adaptation-level as a basis for a quantitative theory of frames of

    reference. *Psychological Review, 55*(6), 297–313. https://doi.org/10.1037/h0056721

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological*

    Bulletin, 114(1), 3-28. doi:10.1037/0033-2909.114.1.3

Jones, E. E. (1990). *Interpersonal perception*. New York: W.H. Freeman.

Kenrick, D. T., McCreath, H. E., Govern, J., King, R., & Bordin, J. (1990). Person-environment

    intersections: Everyday settings and common trait dimensions. *Journal of Personality*

    and Social Psychology, 58(4), 685–698. https://doi.org/10.1037/0022-3514.58.4.685

Koehler, D. J. (1991). Explanation, imagination, and confidence in judgment. *Psychological*

    Bulletin, 110(3), 499–519. https://doi.org/10.1037/0033-2909.110.3.499

Lay, C. H., & Jackson, D.N. (1969). Analysis of generality of trait-inferential relationships. *Journal of Personality and Social Psychology, 12*, 12-21. doi:10.1037/h0027373

Levay, K. E., Freese, J., & Druckman, J. N. (2016). The Demographic and Political Composition of Mechanical Turk Samples. *SAGE Open*, 6(1). https://doi.org/10.1177/2158244016636433

Lord, C. G., & Lepper, M. R. (1999). Attitude representation theory. *Advances in Experimental Social Psychology, 31*, 265-343. doi.org/10.1016/S0065-2601(08)60275-0

Lord, C. G., Paulson, R. M., Sia, T. L., Thomas, J. C., & Lepper, M. R. (2004). Houses built on shifting sand: Effects of exemplar stability on resistance to attitude change. *Journal of Personality and Social Psychology, 87*(6), 733-749. https://doi.org/10.1037/0022-3514.87.6.733

Maio, G. R., & Esses, V. M. (2001). The need for affect: Individual differences in the motivation to approach or avoid emotions. *Journal of Personality, 69*(4), 583–615. https://doi.org/10.1111/1467-6494.694156

McCarty, N. (2019). *Polarization: What everyone needs to know*. New York: Oxford University Press.

Millar, M. G., & Tesser, A. (1986). Thought-induced attitude change: The effects of schema structure and commitment. *Journal of Personality and Social Psychology, 51*(2), 259–269. doi: 10.1037/0022-3514.51.2.259

Montoya, A. K., & Hayes, A. F. (2017). Two-condition within-participant statistical mediation analysis: A path-analytic framework. Psychological Methods, 22(1), 6–27. https://doi.org/10.1037/met0000086

Parsons, K. & Donehoo, P. N. (2019). *Polarized: The collapse of truth, civility, and community in divided times and how we can find common ground*. Prometheus Books.

Petrocelli, J. V., Clarkson, J. J., Tormala, Z. L., & Hendrix, K. L. (2010). Perceiving stability as a means to attitude certainty: The role of implicit theories of attitudes. *Journal of Experimental Social Psychology, 46*, 874-883. https://doi.org/10.1016/J.JESP.2010.07.012

Petty, R. E., & Cacioppo, J. T. (1986). The Elaboration Likelihood Model of persuasion. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 19, pp. 123-205). New York, NY: Academic Press.

Petty, R. E., Briñol, P., & Tormala, Z. L. (2002). Thought confidence as a determinant of persuasion: The self-validation hypothesis. *Journal of Personality and Social Psychology, 82*(5), 722-741. https://doi.org/10.1037/0022-3514.82.5.722

Ramsey, S. L., Lord, C. G., Wallace, D. S., & Pugh, M. A. (1994). The role of subtypes in attitudes towards superordinate social categories. *British Journal of Social Psychology, 33*(4), 387-403. doi:10.1111/j.2044-8309.1994.tb01036.x

Reyna, V. F. (2012). A new intuitionism: Meaning, memory, and development in Fuzzy-trace theory. *Judgment and Decision Making, 7*(3), 332–359.

Reyna, V. F., & Brainerd, C. J. (1995). Fuzzy-trace theory: Some foundational issues. *Learning and Individual Differences, 7*(2), 145–162. doi: 10.1016/1041-6080(95)90028-4

Ross, L., Lepper, M. R., & Hubbard, M. (1975). Perseverance in self-perception and social perception: Biased attributional processes in the debriefing paradigm. *Journal of Personality and Social Psychology, 32*(5), 880–892. https://doi.org/10.1037/0022-3514.32.5.880

Sadler, O., & Tesser, A. (1973). Some effects of salience and time upon interpersonal hostility and attraction during social isolation. *Sociometry, 36*, 99-112. doi: 10.2307/2786285

Schneider, D. J. (1973). Implicit personality theory: A review. *Psychological Bulletin, 79*(5), 294–309. https://doi.org/10.1037/h0034496

Schwarz, N., & Bless, H. (2007). Mental Construal Processes: The Inclusion/Exclusion Model. In D. A. Stapel & J. Suls (Eds.), *Assimilation and contrast in social psychology* (pp. 119–141). Psychology Press.

Schwarz, N., & Bohner, G. (2001). The construction of attitudes. In A. Tesser & N. Schwarz (Eds.), *Blackwell handbook of social psychology, Vol. 1: Intraindividual processes* (pp. 436-457). Oxford, UK: Blackwell.

Schwarz, N., Bless, H., Strack, F., Klumpp, G., Rittenauer-Schatka, H., & Simons, A. (1991). Ease of retrieval as information: Another look at the availability heuristic. *Journal of Personality and Social Psychology, 61*, 195–202. https://doi.org/10.1037/0022-3514.61.2.195

Schwarz, N. (2006). Attitude research: Between Ockham's razor and the fundamental attribution error. *Journal of Consumer Research, 33*(1), 19-21. https://doi.org/10.1086/504124

Schwarz, N. (2007). Attitude construction: Evaluation in context. *Social Cognition, 25*(5), 638-656. doi: 10.1521/soco.2007.25.5.638

Sia, T. L., Lord, C. G., Blessum, K. A., Ratcliff, C. D., & Lepper, M. R. (1997). Is a rose always a rose? The role of social category exemplar change in attitude stability and attitude-behavior consistency. *Journal of Personality and Social Psychology, 72*, 501-514. doi: 10.1037/0022-3514.72.3.501

Sides, J. & Hopkins, D. J. (2015). *Political polarization in American politics*. New York: Bloomsbury Academic.

Slamecka, N. J., & Graf, P. (1978). The generation effect: Delineation of a phenomenon. *Journal of Experimental Psychology: Human Learning and Memory, 4*(6), 592–604. https://doi.org/10.1037/0278-7393.4.6.592

Smith, E. R., & DeCoster, J. (2000). Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review, 4*(2), 108–131. https://doi.org/10.1207/S15327957PSPR0402_01

Spencer, S. J., Zanna, M. P., & Fong, G. T. (2005). Establishing a causal chain: Why experiments are often more effective than mediational analyses in examining psychological processes. *Journal of Personality and Social Psychology, 89*(6), 845–851. https://doi.org/10.1037/0022-3514.89.6.845

Strahan, R., & Gerbasi, K. C. (1972). Short, homogeneous versions of the Marlowe-Crowne Social Desirability Scale. *Journal of Clinical Psychology, 28*(2), 191–193. https://doi.org/10.1002/1097-4679(197204)28:2<191::AID-JCLP2270280220>3.0.CO;2-G

Svedholm-Häkkinen, A. M., & Lindeman, M. (2018). Actively open-minded thinking: Development of a shortened scale and disentangling attitudes towards knowledge and people. *Thinking & Reasoning, 24*(1), 21– 40. https://doi.org/10.1080/13546783.2017.1378723

Tesser, A., & Conlee, M. C. (1975). Some effects of time and thought on attitude polarization. *Journal of Personality and Social Psychology, 31*(2), 262–270. https://doi.org/10.1037/h0076292

Tesser, A., & Cowan, C. L. (1975). Thought and number of cognitions as determinants of attitude change. *Social Behavior and Personality: An International Journal, 3*(2), 165–173. https://doi.org/10.2224/sbp.1975.3.2.165

Tesser, A., & Cowan, C. L. (1977). Some attitudinal and cognitive consequences of thought. *Journal of Research in Personality, 11*(2), 216–226. https://doi.org/10.1016/0092-6566(77)90018-6

Tesser, A., & Leone, C. (1977). Cognitive schemas and thought as determinants of attitude change. *Journal of Experimental Social Psychology, 13*(4), 340–356. https://doi.org/10.1016/0022-1031(77)90004-X

Tesser, A. (1976). Attitude polarization as a function of thought and reality constraints. *Journal of Research in Personality, 10*(2), 183–194. https://doi.org/10.1016/0092-6566(76)90071-4

Tesser, A. (1978). Self-generated attitude change. In L. Berkowitz (Vol. Ed.), *Advances in experimental social psychology. Vol. 11.* (pp. 289–338). New York, NY: Academic Press. doi: 10.1016/S0065-2601(08)60010-6

Tormala, Z.L., Falces, C., Briñol, P., & Petty, R.E. (2007). Ease of retrieval effects in social judgment: The role of unrequested cognitions. *Journal of Personality and Social Psychology, 93*, 143-157. https://doi.org/10.1037/0022-3514.93.2.143

Tormala, Z. L., Petty, R. E., & Briñol, P. (2002). Ease of retrieval effects in persuasion: A self-validation analysis. *Personality and Social Psychology Bulletin, 28*, 1700-1712. https://doi.org/10.1177/014616702237651

Tourangeau, R. (1992). Context effects on response to attitude surveys: Attitudes as memory structure. In N. Schwartz & S. Sudman, (Eds.) *Context Effects in Social and Psychological Research* (pp. 35–48). New York, NY: Springer.

Valenti, A. C. & Tesser, A. (1981). On the mechanism of thought-induced attitude change. *Social Behavior and Personality, 9*(1), 17-22. https://doi.org/10.2224/sbp.1981.9.1.17

Wang, L., Zhang, Z., McArdle, J. J., & Salthouse, T. A. (2008). Investigating ceiling effects in longitudinal data analysis. *Multivariate Behavioral Research, 43*(3), 476-496. https://doi.org/10.1080/00273170802285941

Wason, P. C. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology, 12*(3), 129-140, DOI: 10.1080/17470216008416717

Webster, D. M., Richter, L., & Kruglanski, A. W. (1996). On leaping to conclusions when feeling tired: Mental fatigue effects on impressional primacy. *Journal of Experimental Social Psychology, 32*(2), 181–195. https://doi.org/10.1006/jesp.1996.0009

Wilson, T. D., & Hodges, S. D. (1992). Attitudes as temporary constructions. In L. L. Martin & A. Tesser (Eds*.), The construction of social judgments* (pp. 37-65). Hillsdale, NJ, US: Lawrence Erlbaum Associates, Inc.

Wyer, R. S., Jr., & Srull, T. K. (1989). *Memory and cognition in its social context*. Lawrence Erlbaum Associates, Inc.

Appendix A

Experiment 2 initial and extrapolated traits for each social group with mean trait extremity and mean frequency with which the traits were generated by Experiment 1 participants. Extremity ranges from 0 = *completely neutral* to 5 = *most extreme*. Standard deviations are in parentheses.

Negative Groups				
	Burum Initial	Burum Extrap	Z'dura Initial	Z'dura Extrap
	Fussy	Selfish	Brash	Selfish
	Secretive	Generous	Possessive	Generous
	Critical	Demanding	Overconfident	Arrogant
	Withdrawn	Easygoing	Uncouth	Humble
	Complaining	Annoying	Showy	Loud
	Detached	Agreeable	Boastful	Calm
	Dissatisfied	Rude	Tricky	Rude
	Finicky	Polite	Ostentatious	Polite
M **Extremity**	1.90 (0.25)	3.47 (0.58)	1.93 (0.37)	3.43 (0.71)
Mdn **Frequency**		17.00		31.50
Positive Groups				
	Burum Initial	Burum Extrap	Z'dura Initial	Z'dura Extrap
	Candid	Forgiving	Traditional	Organized
	Congenial	Unforgiving	Methodical	Disorganized
	Gregarious	Kind	Systematic	Honest
	Lenient	Cruel	Cultured	Dishonest
	Casual	Friendly	Prudent	Intelligent
	Obliging	Unfriendly	Refined	Unintelligent
	Pardoning	Arrogant	Pragmatic	Dependable
	Modest	Humble	Straightforward	Undependable
M **Extremity**	2.04 (0.49)	3.76 (0.54)	2.07 (0.54)	3.64 (0.78)
Mdn **Frequency**		21.00		19.50

VITA

Kaleigh Ann Decker is from Terre Haute, Indiana. She attended Indiana State University in Terre Haute, Indiana, where she earned her BS in Psychology with two minors in Sociology and Civic Leadership. After graduating from Indiana State University, Kaleigh attended Texas Christian University (TCU) to study Experimental Social Psychology under the guidance of Dr. Charles Lord. At TCU, Kaleigh earned her MS in 2020 and her PhD in 2022. Kaleigh was hired as a Market Research Strategist at Wonderlust Collective, after completing her degree at TCU.

ABSTRACT

EFFECTS OF BIASED EXTRAPOLATION ON ATTITUDE EXTREMITY

by Kaleigh Ann Decker, M.S., 2020

Department of Psychology

Texas Christian University

Dissertation Advisor: Dr. Charles Lord, Professor of Psychology

Past research has found that thinking about an attitude object can make positive and negative attitudes more extreme. The current research explored whether a certain type of thought—*extrapolating* from known to unknown values—would make attitudes and behavioral intentions more extreme than *reviewing* known values. In three experiments, I found that extrapolating from a social groups known personality traits made positive and negative attitudes and behavioral intentions more extreme than reviewing a social group's known personality traits. This pattern of results occurred when participants self-generated their own extrapolations (Experiment 1) or rated the likelihood of frequently generated additional traits (Experiment 2). Attitudes were also more extreme after extrapolating to traits high versus low in cognitive relevance to the social group's known traits, regardless of the positivity or negativity of the extrapolated traits (Experiment 3). In all three experiments, the effect of biased trait extrapolation on attitude extremity was mediated by more extreme associations to the extrapolated social group. The current findings are consistent with attitude construal theories, which suggest attitudes are informed by accessible associations.