



Perceiving attitude change: How qualitative shifts augment change perception

Christopher J. Bechler^{a,*}, Zakary L. Tormala^a, Derek D. Rucker^b

^a Stanford University, United States of America

^b Northwestern University, United States of America



ARTICLE INFO

Handling editor: Tara MacDonald

Keywords:

Attitudes

Attitude change

Perceived change

Qualitative change

Persuasion

ABSTRACT

Attitude change and persuasion are among the most studied topics in social psychology. Surprisingly, though, as a field we have virtually zero insight into *perceived* attitude change—that is, how people assess the magnitude of a shift in someone's attitude or opinion. The current research provides an initial investigation of this issue. Across 6 primary experiments and a series of supplemental studies (total $N = 2880$), we find consistent support for a *qualitative change hypothesis*, whereby qualitative attitude change (change of valence; e.g., from negative to positive) is perceived as greater than otherwise equivalent non-qualitative attitude change (change within valence; e.g., from negative to less negative or from positive to more positive). This effect is mediated by ease of processing: Qualitative attitude change is easier for people to detect and understand than non-qualitative attitude change, and this ease amplifies the degree of perceived change. We examine downstream consequences of this effect and discuss theoretical, methodological, and practical implications.

Imagine a public health campaign designed to change parents' attitudes toward vaccinating their children. The campaign is pro-vaccine and is delivered to three groups of parents. One group is strongly opposed to vaccines, another group is moderately opposed, and a third group is moderately in favor. On a scale ranging from 1 (strongly opposed) to 9 (strongly in favor), with a neutral midpoint of 5, the average of the first group is a 2, the second group a 4, and the third group a 6. Now imagine that individuals in all of these groups receive the same message endorsing vaccines, and on average it moves each group two points. The first group is now just moderately opposed (4), the second group moderately in favor (6), and the third group strongly in favor (8). Which group would campaign officials perceive to have changed the most? In which case would observers conclude the message was most persuasive?

Attitude change and persuasion are among the most highly studied topics in social psychology. Decades of research, and hundreds of articles, have been devoted to understanding them. This literature has produced a litany of insights into the factors that shape people's attitudes and cause them to become more or less favorable toward a vast array of objects and issues (for reviews see Petty & Wegener, 1998; Tormala & Briñol, 2015). Interestingly, though, despite the field's collective interest in these topics, surprisingly little is known about the factors that drive *perceived* attitude change. That is, what shapes people's assessments of how much attitude change has occurred in a given

individual or situation? People witness attitude change on a daily basis—learning, for example, about a friend's souring on a new restaurant or a politician's change of course on a social issue—but we know little to nothing about the determinants of whether this change seems large or small. The current research examines this issue. Specifically, we examine whether and how the *type* of attitude change someone experiences (e.g., from very negative to less negative, from negative to positive, or from somewhat positive to very positive) affects the perceived magnitude of that change—that is, the extent to which people believe a large or small amount of change has occurred. If it does, we submit that differences in the perceived amount of change that has occurred can affect the perceived importance or significance of that change for behavioral outcomes.

1. Perceived attitude change

Although we found no research examining the factors that make attitude change seem large or small, a literature does exist on perceived attitude change more generally. For example, some prior work has uncovered asymmetries in people's perceptions of their own versus others' attitude change. Whereas people are often inaccurate at estimating their own change over time (e.g., Ross, 1989), they appear more accurate at estimating others' change (Douglas & Sutton, 2004). Likewise, people think negative media messages impact others more than

* Corresponding author.

E-mail address: cbechler@stanford.edu (C.J. Bechler).

themselves (Davison, 1983), but see positive messages as affecting themselves more than others (e.g., Duck, Terry, & Hogg, 1995). We also know that people's perceptions of their own change (versus stability) can have implications for attitude certainty (e.g., Petrocelli, Clarkson, Tormala, & Hendrix, 2010; Tormala & Petty, 2002), and that the perceived malleability of others' attitudes can affect people's willingness to try to persuade them (Akhtar & Wheeler, 2016). In addition, research suggests that people form positive and negative attributions about others when they observe those others change their attitudes or attitude-relevant behaviors (e.g., Barden, Rucker, & Petty, 2005; Klein & O'Brien, 2017; Kreps, Laurin, & Merritt, 2017; Reich & Tormala, 2013).

Importantly, though, none of this prior work asks, let alone answers, the question of interest in the current research: What shapes people's perceptions that the amount of attitude change someone has experienced is large or small? We explore the possibility that a crucial factor that affects perceptions of change magnitude is the *type* of attitude change under consideration—that is, the point along an evaluative continuum at which change occurs. Consider the opening example of the vaccine campaign. One group changed from strongly opposed to just moderately opposed, another from moderately opposed to moderately in favor, and another from moderately in favor to strongly in favor. Would this difference in the type, or location, of attitude change affect people's perceptions of how much change occurred? On the face of it, multiple possibilities exist.

1.1. Null hypothesis

One possibility, the null hypothesis, is that the type or location of change that has occurred has no impact on the perceived magnitude of change. Especially in contexts in which people can easily calculate the degree of attitude change that has taken place (e.g., a survey demonstrates that an intervention produced a 2-point shift on a 9-point scale), and in which that calculation is highly salient, it could be that people attend solely to the raw, or absolute, magnitude of change to assess perceived change. In other words, people might be sensitive to the objective amount of change that has occurred and insensitive to the type of change involved. If true, in the vaccine example we would expect people to perceive an equivalent amount of change for each group of message recipients. Far from a strawman, this outcome would appear to reflect a logical approach to the question: If people know that everyone changed 2 points on a scale, they could reasonably view that change as equivalent regardless of where on the scale it occurs.

In fact, this perception would be consistent with the way attitudes researchers treat attitude change. In the attitudes literature, a common method to assess change is to measure attitudes both before and after some treatment or manipulation and compare differences in the degree of change from Time 1 to Time 2. In persuasion studies, for example, researchers often measure attitudes before and after a persuasive message or activity, manipulate some aspect of that message or activity, and then compare the degree of change across conditions (e.g., Blankenship, Wegener, & Murray, 2012; Clarkson, Tormala, & Leone, 2011; Janis & King, 1954; Tormala & Petty, 2002). The assumption is that wherever the pre-post difference is greater, more change has occurred. Likewise, if the pre-post difference is similar across conditions, it is assumed that attitude change is equivalent. To our knowledge, there has been no indication that some pre-post differences are viewed as greater simply by virtue of their location along an evaluative continuum. Two-point change is treated as two-point change, in other words, regardless of where it occurs on an attitude scale. The lack of attention to this issue suggests that attitude researchers, at least implicitly, operate in a manner consistent with the null hypothesis.

1.2. Extremity hypotheses

An alternative hypothesis is that the type or location of change does affect people's perceptions of change. An attitude strength perspective

(Petty & Krosnick, 1995), in particular, might predict that people perceive greater change when that change involves a shift in attitude extremity. Attitude extremity is a central dimension of attitude strength (e.g., Abelson, 1995; Bassili, 1996; Krosnick, Boninger, Chuang, Berent, & Carnot, 1993) and strong attitudes are well-known to be more resistant to change and more influential over behavior than weak attitudes (e.g., Krosnick & Petty, 1995). If laypeople intuit that extreme attitudes are stronger than moderate attitudes, they might believe it would take greater effort to change a person who has a more extreme initial attitude. If true, people might infer greater change and/or a more effective persuasive message when someone moves from an extreme starting point to a moderate ending point. In the vaccine example, this inference might elicit greater perceived change for the first group, who moved from strongly opposed (extreme) to moderately opposed (moderate).

On the other hand, an attitude strength perspective could predict the opposite—that shifts *from moderate to extreme* constitute the greatest change, because ending with an extreme attitude is most likely to trigger behavior or prompt people to take action. Again, strong attitudes are more influential over behavior, and extremity is a central dimension of attitude strength. Perhaps people intuit this relationship, anchor on the final attitude, and conclude that the most change has occurred when people end up the most extreme. In the vaccine example, the third group, who moved from moderately in favor to strongly in favor, might be perceived as most likely to take action and vaccinate their children, fostering the inference that this group changed the most.

In short, if the type or location of attitude change affects perceived change at all, people might perceive greater change when that change results in less extreme or more extreme attitudes—the first and third groups in the vaccine example.

2. The qualitative change hypothesis

Although the null and extremity hypotheses have merit, we posit that people generally perceive greater attitude change when that change crosses a neutral threshold and shifts from one valence to another. In essence, we propose that people perceive qualitative change as greater than non-qualitative change. In the attitudes context, this means that changes of valence (or changes of type; e.g., from negative to positive) would be viewed as greater than changes *within* valence (or changes of degree; e.g., from very negative to less negative or from somewhat positive to more positive). Operationally, we submit that even if the absolute magnitude of attitude change on a scale is held constant—say, 2 points on a 9-point attitude scale—people will perceive greater change when that change shifts valence, or crosses the neutral midpoint, compared to when it does not. We refer to this prediction as the *qualitative change hypothesis*, postulating that qualitative change (i.e., change of valence) will tend to be viewed as greater than mathematically identical non-qualitative change (i.e., change *within* valence).¹

Why would qualitative attitude change be perceived as greater than non-qualitative attitude change? We draw from prior research on processing ease, or fluency, and propose that qualitative change is viewed as greater than non-qualitative change because it is easier for people to detect and understand. Indeed, an extensive literature suggests that people often find it easier to distinguish between stimuli or states that

¹ Note that although the present research focuses on changes of valence—that is, on change that crosses a neutral threshold—qualitative change might be seen as greater than non-qualitative change in other contexts as well (e.g., when an attitude changes in basis; from cognitive to emotional). In this initial research we focus on valence and neutral thresholds given their centrality to the attitude construct. We consider other forms of qualitative change, and outline directions for future research on this topic, in the [General discussion](#).

differ qualitatively, or categorically, rather than simply quantitatively. Furthermore, past research reveals that people tend to evaluate objects and entities as more different from each other when a categorical boundary—even an arbitrary one—exists between them. The current research merges these insights to illuminate the effect of qualitative versus non-qualitative evaluative shifts on perceived attitude change.

2.1. Qualitative differences are easier to process

First, considerable research suggests that equivalent differences between pairs of stimuli can be easier to detect when those differences cross a boundary of some kind. For example, research on categorical perception has revealed that people are better able to discriminate between things “when those things belong to different categories rather than the same category, controlling for the physical difference between the things” (Goldstone & Hendrickson, 2010, p. 69; see also Harnad, 1987; Pastore, 1987). As one illustration, Bornstein and Korda (1984) found that participants could discriminate more quickly between colors of different hues (i.e., colors that differed qualitatively) than colors of the same hue (i.e., colors that did not differ qualitatively), even when the difference in wavelength between the colors—that is, the physical distance between stimuli—was held constant. In addition to being quicker, cross-category color discrimination is more accurate and reliable than within-category color discrimination (e.g. Bornstein, 2006; Bornstein & Korda, 1984; Roberson & Davidoff, 2000). Similar effects have been observed in research using sounds and phonemes, faces, and even emotions as stimuli (e.g., Beale & Keil, 1995; Etcoff & Magee, 1992; Liberman, Harris, Hoffman, & Griffith, 1957). Assuming response speed and accuracy are at least somewhat correlated with the experience of processing ease (e.g., see Attneave, 1955; Reber, Winkielman, & Schwarz, 1998), these findings hint at the possibility that people might have an easier time perceiving attitude change when that change creates categorically different Time 1 and Time 2 attitudes, such as attitudes of different valence.

Also germane, research on processing fluency suggests that people experience greater ease of processing when the stimuli to be processed have high rather than low contrast (Reber et al., 1998; Jacoby, Kelley, Brown, & Jasechko, 1989; see Alter & Oppenheimer, 2009). For example, people find it easier to read information when it is presented in highly contrasting color combinations, which can have a host of consequences including making the information seem truer, more legitimate, and more persuasive (e.g., Briñol, Petty, & Tormala, 2006; Reber & Schwarz, 1999; Whittlesea, Jacoby, & Girard, 1990). To the extent that qualitative attitude change can make initial and subsequent attitudes appear more contrasting, these findings lend credence to the notion that attitude change might be easier to process when it involves a qualitative rather than non-qualitative shift.

2.2. Qualitative differences appear larger

If qualitative change is easier to process, people might readily perceive it as *greater*. Consistent with this notion, differences between stimuli do tend to be perceived as larger in magnitude when they span a categorical boundary of some kind. In some of the earliest work on categorical judgment, Tajfel and Wilkes (1963) proposed that classifying, or categorizing, stimuli produces an increase “in the apparent differences between stimuli belonging to different classes, and in the apparent similarity of stimuli belonging to the same class” (p. 101). Assuming people classify positive and negative attitudes as different, this finding suggests that people might perceive movement from negative to positive, or vice versa, as greater—that is, as spanning a greater distance—than movement within negative or positive.

Recent support for this proposition comes from research on boundary effects in judgment and decision making. Lacetera, Pope, and Sydnor (2012) found that differences in sale prices between used cars are greater when those cars' mileage differences cross a 10,000-mile

boundary (e.g., 59,000 versus 61,000) compared to when they do not (e.g., 61,000 versus 63,000). Krueger and Clement (1994) found that people estimate that two days in adjacent months differ more in temperature than two equivalently-spaced days within the same month. Mishra and Mishra (2010) showed that the presence (versus absence) of a state border between a city and a disaster site reduces the perceived risk associated with a disaster, even when the actual distance between the city and disaster site is constant. Isaac and Schindler (2013) documented that rank improvements that cross category boundaries (e.g., moving from 11th to 10th ranked, now in the “top-ten”) seem greater than equivalent improvements that do not cross category boundaries (e.g., moving from 10th to 9th). In short, considerable evidence suggests that identical changes, or differences, can seem greater when they span a categorical boundary.

2.3. Ease and magnitude are inherently linked

In sum, research from a variety of domains suggests that categorical differences between stimuli are easier to process and thus perceived as greater than equivalent differences within category. Extending these insights to attitudes, we postulate that qualitative attitude change (e.g., from negative to positive) is perceived as greater than non-qualitative change (e.g., from very to less negative or from somewhat to more positive), because qualitative change is easier to detect and understand and this ease fosters an inference of greater magnitude. In addition to being grounded in past research, this prediction is rooted in everyday human experience. People find it easier to perceive, or process, stimuli that are larger in magnitude. For instance, it is easier for most people to read large rather than small font, to see something in the distance (e.g., a mountain) if it is large as opposed to small, and to hear a loud (high decibel) as opposed to soft (low decibel) sound. Thus, ease and magnitude are inherently linked (see also Thomas & Morwitz, 2009). We harness this insight to shed new light on the perception of qualitative versus non-qualitative attitude change.

3. Implications of the qualitative change hypothesis

The qualitative change hypothesis suggests that qualitative attitude change is easier than non-qualitative attitude change for people to detect and understand, and this ease increases its perceived magnitude. If correct, this hypothesis has numerous implications. First, at a general conceptual level, it would offer initial insight into an unanswered question in the attitudes literature: What shapes people's perceptions that attitude change is large or small? As noted, several viable hypotheses could be derived from extant literature. Evidence for the qualitative change hypothesis would thus yield novel insight with potential implications for how we understand evaluative shifts in a variety of contexts.

Second, if correct, the qualitative change hypothesis could have important interpersonal implications. For instance, it could affect how individuals evaluate others. Consider a political campaign in which opposing candidates all shift their opinions on a particular issue, such as gun control. The qualitative change hypothesis suggests that if one of those candidates changes the valence of her view and the others merely change in extremity (becoming less or more extreme than they were previously), the one who changes in valence could be seen as having shifted more since the campaign's inception. If true, this perception could affect people's willingness to vote for her in an election (e.g., Kreps et al., 2017; Reich & Tormala, 2013).

Finally, the qualitative change hypothesis could have important methodological implications. As noted, attitudes researchers generally treat Time 2 – Time 1 attitude change as a monolithic construct, typically paying no attention to whether that change occurs within or across valence. This tendency reflects a tacit assumption that all change can be treated equally, regardless of its location on an evaluative continuum. The qualitative change hypothesis suggests that this

assumption might warrant reassessment. If attitude change is perceived as greater when it crosses a neutral midpoint, seemingly similar change across experimental conditions could mask differences in people's perceptions of that change, highlighting a potentially important disconnect between how researchers and laypeople interpret evaluative shifts.

4. Overview

We present six experiments testing the effect of attitude change type—qualitative versus non-qualitative—on perceived attitude change. Experiments 1–4 develop a paradigm to provide an initial test of the qualitative change hypothesis and assess its robustness to procedural variations. Experiments 5 and 6 examine the psychological mechanism and consequences of the qualitative change effect, respectively. Across experiments, we make an effort to have sufficient statistical power by collecting data from at least 100 participants per experimental condition (200 per condition when testing mechanism). All manipulations and conditions are reported in the main text, and all measures and results are reported in the main text or Online supplement. We also report any participant exclusions.

5. Experiment 1

In Experiment 1, participants read about a scenario in which one person tried to persuade another person to change his attitude toward a policy. All participants read that the target of the persuasion attempt did, in fact, change his attitude. Importantly, though, the nature of the target's attitude change was varied across three conditions. In one—the qualitative change condition—participants read that target changed his attitude from negative to positive on the focal issue. In the other conditions—the non-qualitative change conditions—participants read that the target changed his attitude, but this change occurred within valence. In one case, he changed from very negative to less negative on the focal issue; in the other, he changed from somewhat positive to more positive on the issue. We predicted that participants would perceive qualitative change as significantly greater than non-qualitative change.

To provide a clean test of the qualitative change hypothesis, we sought to hold the direction and approximate amount of change constant across conditions. However, in this initial experiment, we also wanted to vary the target's attitude change using verbal descriptions absent any information about the target's attitude on an attitude scale (e.g., without noting that he was a 6 on a 1–9 scale). The goal was to offer a general demonstration of the qualitative change effect without capitalizing on potential ambiguity in participants' perceptions of attitude scales (which typically are ordinal in nature) and how much distance separates individual scale points (see Eagly & Chaiken, 1993; Stevens, 1946).

5.1. Method

5.1.1. Participants and design

A total of 304 participants, recruited through Amazon's Mechanical Turk, took part for monetary compensation. Participants were randomly assigned to one of three change conditions: negative to less negative, negative to positive (qualitative change), or positive to more positive. A sensitivity analysis (GPower; Faul, Erdfelder, Lang, & Buchner, 2007) revealed that our sample provided 80% power to detect an effect size of $\eta^2 = 0.03$ ($d = 0.35$).

5.1.2. Procedure

All participants read about a scenario in which one person, Bob, attempted to change another person's, Frank's, mind about a policy. They read that Bob supported the policy, made some interesting points, and ultimately moved Frank's opinion. After participants read the scenario, they completed a series of measures assessing perceived change.

5.1.3. Change condition

Participants were randomly assigned to one of three change conditions. In the negative to positive (i.e., qualitative change) condition, participants read that Frank's attitude changed in valence after hearing Bob's message. Specifically, participants read that Frank was “somewhat opposed to the policy” at the beginning of the interaction, but that he shifted and was now “somewhat in favor” after hearing Bob's message. The scenarios for the non-qualitative change conditions were similar, except that Frank's change was described as happening within valence. In the negative to less negative condition, participants read that Frank was “pretty opposed to the policy” at the beginning of the interaction, but that he shifted and was now “just somewhat opposed” after hearing Bob's message. In the positive to more positive condition, participants read that Frank was “just somewhat in favor of the policy” at the beginning of the interaction, but that he shifted and was now “pretty in favor” after hearing Bob's message. Our aim was to keep the direction and approximate amount of change constant across conditions, modifying only the type of change described.

5.1.4. Dependent measures

5.1.4.1. Perceived change. Our primary dependent measure was a composite index of four items assessing perceived change. Two of the items assessed perceptions of the change caused by Bob and two of the items assessed perceptions of Frank's change. For the former, participants rated how much change Bob caused in Frank's views and how persuasive Bob was. For the latter, participants rated how persuaded Frank was and how much Frank changed his opinion. Participants responded to each item on a 1–9 scale, scored such that higher values reflected perceptions of greater change. Responses to these measures were averaged to form a composite index ($\alpha = 0.90$). In addition, participants answered eight exploratory questions about Bob and Frank. These measures and their results are detailed in the Online supplement.

5.1.4.2. Time 1 and Time 2 attitude estimates. Because we relied on descriptive labels (e.g., “somewhat opposed” and “somewhat in favor”) to describe Frank's initial and final attitude, it was important to assess whether we accurately calibrated Frank's attitude positions. That is, it was important to establish that Frank's starting and ending positions were approximately equidistant from each other across conditions, which would make for a clear comparison between qualitative and non-qualitative change. Thus, we asked participants to estimate Frank's initial and final attitudes (i.e., his positions before and after hearing Bob's message) on a scale ranging from 0 to 100, where 0 meant “strongly opposed,” 100 meant “strongly in favor,” and 50 meant “neutral.” We then calculated the difference between those estimates for each condition.

5.2. Results

Two participants failed to complete all of the perceived change measures and were excluded from analysis. For the remaining 302 participants, we submitted the perceived change index and Time 1 and Time 2 attitude estimates (and their difference) to a one-way analysis of variance (ANOVA) with change condition as the independent variable.

5.2.1. Perceptions of change

First, we found a significant effect of the manipulation on perceived change, $F(2, 299) = 13.92$, $p < .001$, $\eta^2 = 0.09$ ($d = 0.63$; see Table 1). Consistent with the qualitative change hypothesis, planned comparisons revealed that change from negative to positive was perceived as greater than both change from negative to less negative, $t(299) = -5.27$, $p < .001$, $d = -0.72$, and change from positive to more positive, $t(299) = -2.31$, $p = .02$, $d = -0.33$. Though not relevant to our theorizing, these latter two conditions also differed from each other: Participants perceived greater change when Frank moved

Table 1
Perceived change and attitude estimates in Experiment 1. Subscripts should be interpreted within rows only. Values with the same subscript do not differ from each other at a $p < .05$ level.

	Change condition		
	Negative to less negative	Negative to positive	Positive to more positive
Dependent variables as a function of change condition in Experiment 1			
Change composite			
<i>M</i>	5.96 _a	7.05 _b	6.58 _c
<i>SD</i>	1.49	1.51	1.35
<i>N</i>	96	102	104
Frank's attitude (<i>M</i>)			
Frank's Time 2 attitude	45.18 _a	61.81 _b	76.50 _c
Frank's Time 1 attitude	17.95 _a	31.04 _b	49.37 _c
Difference	27.23 _a	30.77 _a	27.13 _a

from positive to more positive rather than from negative to less negative, $t(299) = -3.01$, $p = .003$, $d = -0.43$.

5.2.2. Time 1 and Time 2 attitude estimates

To assess whether we accurately calibrated the attitude position descriptions and described change of equal magnitude (objectively) in each condition, we analyzed the Time 1 and Time 2 attitude estimates and their difference (see Table 1). As expected for both Time 1 and Time 2 attitude estimates, we found a significant effect across conditions ($F_{time 1}(2, 299) = 55.44$, $p < .001$, $\eta^2 = 0.27$, $d = 1.22$; $F_{time 2}(2, 299) = 40.90$, $p < .001$, $\eta^2 = 0.21$, $d = 1.03$). Participants estimated Frank's initial and final attitudes to be more positive in the positive to more positive condition than in the negative to positive and negative to less negative conditions ($F_{time 1}(1, 299) = 92.07$, $p < .001$, $\eta^2 = 0.22$, $d = 1.06$; $F_{time 2}(1, 299) = 58.94$, $p < .001$, $\eta^2 = 0.15$, $d = 0.84$), which also differed from each other ($F_{time 1}(1, 299) = 18.80$, $p < .001$, $\eta^2 = 0.05$, $d = 0.46$; $F_{time 2}(1, 299) = 22.86$, $p < .001$, $\eta^2 = 0.06$, $d = 0.51$). As intended, however, we found no differences in Time 2 – Time 1 difference scores, $F(2, 299) = 0.50$, $p = .61$, $\eta^2 = 0.003$ ($d = 0.11$), suggesting that our attitude position labels accurately calibrated Frank's absolute degree of change across conditions.²

5.3. Discussion

Experiment 1 provided initial evidence for the qualitative change hypothesis: Participants perceived change of valence (qualitative change) as greater than equivalent change within valence (non-qualitative change). We demonstrated this effect without pinpointing, or specifying, the target's Time 1 and Time 2 attitudes on an attitude scale, suggesting that the effect is not a function of how people perceive attitude scales per se, but rather stems from perceptions of change at different points along a general evaluative continuum. It is useful, however, to demonstrate the effect while holding the degree of change perfectly constant across conditions. Despite our effort to achieve this level of control in Experiment 1, there was some ambiguity (i.e., error)

²Notably, in the positive to more positive condition, participants' average estimate of Frank's initial attitude position was basically neutral (49.37). As it turns out, a number of participants made inaccurate estimates (e.g., rating Frank's attitude as slightly negative when we described it as positive). Because we did not plan to exclude participants based on these errors, we made the conservative choice to include all participants in the main analyses. Excluding participants who made errors in either estimate (remaining $n = 180$), we still found a significant effect of change condition on perceived change. Participants perceived greater change in the qualitative rather than non-qualitative change conditions, $F(2, 177) = 33.91$, $p < .001$, $\eta^2 = 0.28$ ($d = 1.25$).

in participants' estimates of Frank's initial and final attitude positions (see Table 1). Specifically, in the positive to more positive condition, participants' mean Time 1 attitude estimates hovered at or just below the neutral point of the scale, rather than above it, as intended. To address this issue going forward, and hold constant the absolute magnitude of the target person's change, we explicitly informed participants of the target's numerical attitude positions and the number of points moved in subsequent experiments. However, we varied the attitude scales employed and the specific attitude positions described to demonstrate the effect across numerous scale types and attitude locations.

Experiment 1 also revealed an unanticipated effect whereby perceived change was greater in the positive to more positive condition than in the negative to less negative condition. As it turns out, several of the subsequent experiments also show this asymmetry—in general, change that happens within positive valence appears to be judged as greater than change that happens within negative valence, regardless of whether it involves change to become more positive or less positive. This effect was not anticipated, nor is it relevant to our primary hypothesis, but we discuss it further in the [General discussion](#).

6. Experiment 2

Experiment 2 tested the qualitative change hypothesis in a context in which a target's movement on an attitude scale was explicitly held constant. To do so, we used a similar scenario as in Experiment 1, but added information about the precise location and distance of Frank's movement on an attitude scale. Based on the qualitative change hypothesis, we postulated that even when the target's attitude change was quantified and identical across conditions, participants would perceive qualitative change as greater than non-qualitative change.

6.1. Method

6.1.1. Participants and design

A total of 296 participants, from Amazon's Mechanical Turk, took part for monetary compensation. Participants were randomly assigned to one of three change conditions: negative to less negative, negative to positive, or positive to more positive. A sensitivity analysis revealed 80% power to detect a minimum effect size of $\eta^2 = 0.03$ ($d = 0.35$).

6.1.2. Procedure

As in Experiment 1, all participants read about a scenario in which one person, Bob, attempted to change another person's, Frank's, mind about a particular policy. They read that Bob made interesting points and moved Frank's opinion. This time, however, participants received information about Frank's initial and final attitudes (before and after hearing Bob's message) on a scale ranging from 0 to 100, where 0 meant “strongly opposed,” 100 meant “strongly in favor,” and 50 meant “neutral.” Notably, 101-point attitude scales have been used with some frequency in past research (e.g., [Abelson, Kinder, Peters, & Fiske, 1982](#); [Greenwald, McGhee, & Schwartz, 1998](#); [Haddock, Zanna, & Esses, 1993](#); [Karpinski & Hilton, 2001](#)). All participants were explicitly informed that “Frank's position moved 16 points up” on this scale, after which they completed measures of perceived change.

6.1.3. Change condition

Participants were randomly assigned to one of three change conditions and received a verbal description of Frank's attitude change. Unlike Experiment 1, however, participants were informed of the specific position of Frank's initial and final attitude. In the negative to positive condition, participants read that Frank was “somewhat opposed to the policy” at the beginning of the interaction, but that he shifted and was now “somewhat in favor” after hearing Bob's message, moving from 47 to 63 on the attitude scale. In the negative to less negative condition, participants read that Frank was “pretty opposed to the policy” at the beginning of the interaction, but that he shifted and

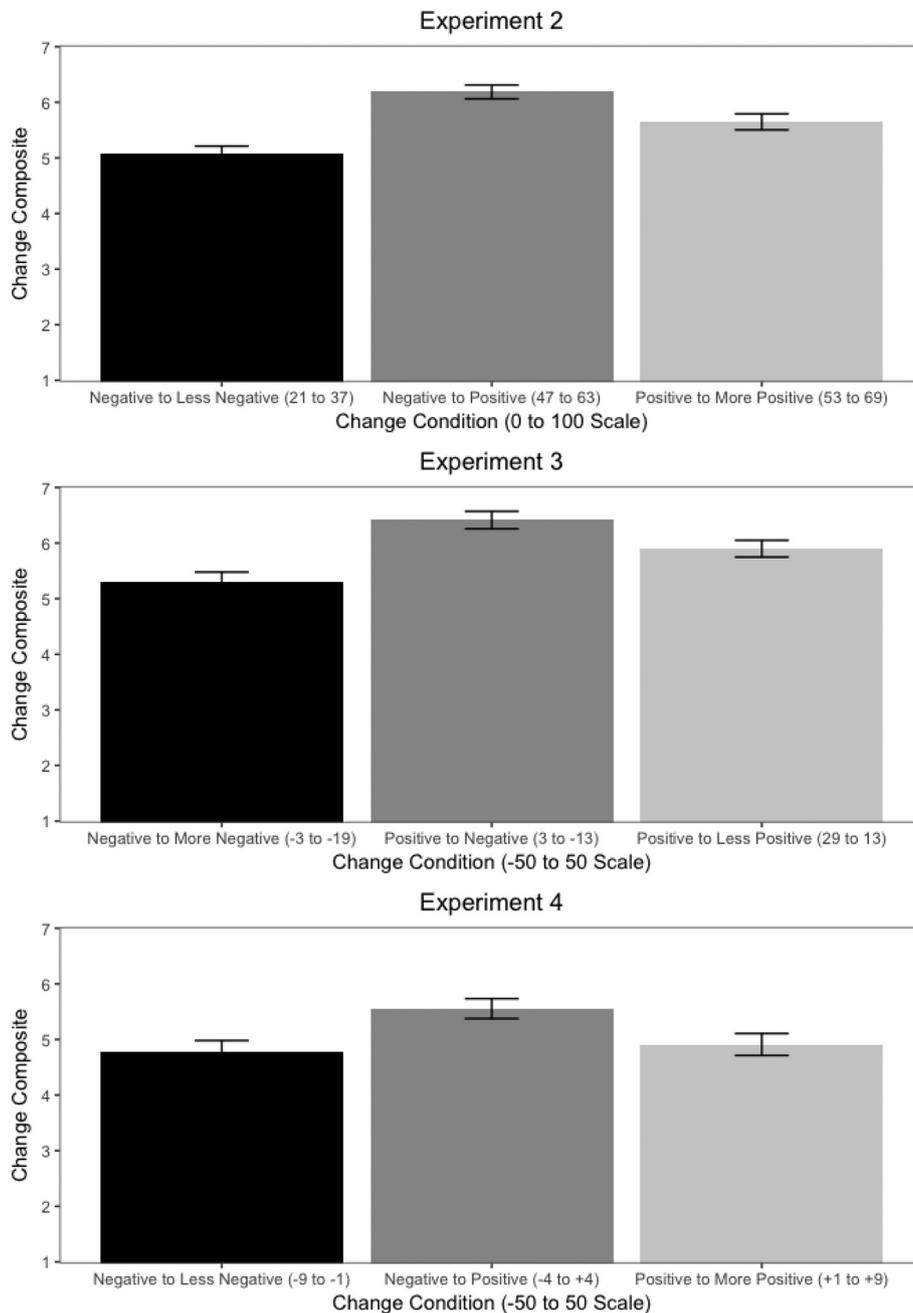


Fig. 1. Perceptions of change in Experiments 2 through 4. Error bars indicate standard errors.

was now “less opposed,” moving from 21 to 37 on the attitude scale. In the positive to more positive condition, participants read that Frank was “somewhat in favor of the policy” at the beginning of the interaction, but that he shifted and was now “even more in favor,” moving from 53 to 69 on the attitude scale. Again, in each condition we explicitly noted that Frank’s attitude moved 16 points on the attitude scale.

6.1.4. Perceived change

Our perceived change measure was the same index of four items employed in Experiment 1 ($\alpha = 0.88$). In addition, participants answered nine exploratory questions (see the Online supplement).

6.2. Results and discussion

We found a significant effect of change condition on perceived change, $F(2, 293) = 16.95, p < .001, \eta^2 = 0.10$ ($d = 0.67$; see Fig. 1).

Planned comparisons indicated that change from negative to positive ($M = 6.19; SD = 1.23$) was perceived as significantly greater than both change from negative to less negative ($M = 5.08; SD = 1.31$), $t(293) = -5.82, p < .001, d = -0.87$, and change from positive to more positive ($M = 5.65; SD = 1.44$), $t(293) = -2.85, p = .005, d = -0.40$. Again, though not germane to our primary concerns, participants perceived that Frank changed more when he moved from positive to more positive rather than from negative to less negative, $t(293) = -3.00, p = .003, d = -0.41$. In short, Experiment 2 replicated the key result from Experiment 1, despite explicitly holding constant the absolute amount of change exhibited.

7. Experiment 3

Experiment 3 further tested the robustness of the qualitative change effect by making several procedural changes. We used a similar

scenario as in Experiments 1 and 2, but changed the valence of Bob's attitude, the direction of Frank's movement, and the attitude scale itself. In this experiment Bob sought to change Frank's attitude by moving it in the *negative* direction, and we described Frank's attitude along a -50 to 50 scale. Because the qualitative change hypothesis is about crossing a threshold—a valence threshold in the present research—and is not tied to the starting or ending valence per se, we expected participants to continue to perceive qualitative change as greater than non-qualitative change.

7.1. Method

7.1.1. Participants and design

A total of 302 participants from Amazon's Mechanical Turk took part for monetary compensation. Participants were randomly assigned to one of three change conditions: negative to more negative, positive to negative, or positive to less positive. A sensitivity analysis revealed 80% power to detect a minimum effect size of $\eta^2 = 0.03$ ($d = 0.35$).

7.1.2. Procedure

As in Experiments 1 and 2, participants read about a scenario in which Bob attempted to change Frank's mind about a policy. However, here they read that Bob was *against* the policy, made interesting points, and moved Frank's opinion in the negative direction. Participants received information about Frank's initial and final attitudes on a scale ranging from -50 to 50 , where -50 meant "strongly opposed," 50 meant "strongly in favor," and 0 meant "neutral," and all participants were explicitly informed that "Frank's position moved 16 points down" on this scale. Following the scenario, participants reported perceived change.

7.1.3. Change condition

In the positive to negative condition, participants read that Frank was "somewhat in favor of the policy" at the beginning of the interaction, but that he shifted and was now "somewhat opposed" after hearing Bob's message, moving from 3 to -13 on the attitude scale. In the negative to more negative condition, participants read that Frank was "somewhat opposed to the policy" at the beginning, but that he shifted and was now "even more opposed," moving from -3 to -19 on the attitude scale. In the positive to less positive condition, participants read that Frank was "pretty in favor of the policy" at the beginning, but that he shifted and was now "less in favor," moving from 29 to 13 on the attitude scale. In each condition, we explicitly noted that Frank's attitude moved 16 points down on the scale.

7.1.4. Perceived change

Our primary measure was the same 4-item index as in Experiments 1 and 2 ($\alpha = 0.88$). Participants also answered six exploratory questions (see Online supplement).

7.2. Results and discussion

We found a significant effect of change condition on perceived change, $F(2, 299) = 11.52$, $p < .001$, $\eta^2 = 0.07$ ($d = 0.55$; see Fig. 1). As hypothesized, change from positive to negative ($M = 6.41$; $SD = 1.56$) was perceived as greater than change from negative to more negative ($M = 5.29$; $SD = 1.91$), $t(299) = -4.79$, $p < .001$, $d = -0.64$, and change from positive to less positive ($M = 5.90$; $SD = 1.50$), $t(299) = -2.17$, $p = .03$, $d = -0.34$. Again, though not relevant to our primary concerns, these latter two conditions also significantly differed from each other: Participants perceived that Frank changed more when he moved from positive to less positive rather than from negative to more negative, $t(299) = -2.59$, $p = .01$, $d = -0.35$. Most germane to our primary concerns, qualitative change was again perceived as greater than non-qualitative change, despite using a different attitude scale and describing a different direction of change.

8. Experiment 4

Experiment 4 further tested the robustness of the qualitative change effect. We adapted the scenario from Experiments 1–3 to change the magnitude and position of Frank's change as well as the way in which Frank's change was described. In Experiments 2–3, we manipulated change somewhat asymmetrically in that the qualitative change condition described change that was more similar to that described in the condition in which Frank became more (as opposed to less) extreme. The goal was to create final attitudes of similar extremity in at least two of the conditions to reduce the number of varying elements. Experiment 4 eliminated this asymmetry and manipulated Frank's change symmetrically around the midpoint of the attitude scale.

8.1. Method

8.1.1. Participants and design

A total of 302 participants from Amazon's Mechanical Turk took part for monetary compensation. Participants were randomly assigned to one of three change conditions: negative to less negative, negative to positive, or positive to more positive. A sensitivity analysis revealed 80% power to detect a minimum effect size of $\eta^2 = 0.03$ ($d = 0.35$).

8.1.2. Procedure

Participants again read the Bob and Frank scenario. As in Experiment 3, all participants received information about Frank's initial and final attitudes on a scale ranging from -50 to 50 , where -50 meant "strongly opposed," 50 meant "strongly in favor," and 0 meant "neutral." This time, however, Frank moved in the positive direction, described his own change, and changed less than in Experiment 3. Following the scenario, they completed measures of perceived change.

8.1.3. Change condition

In the negative to positive condition, participants read that Frank reported he was originally "somewhere around a -4 " on the attitude scale but that he shifted to "somewhere around a $+4$ " after hearing Bob's message. In the negative to less negative condition, Frank reported he was originally "somewhere around a -9 " but shifted to "somewhere around a -1 " after hearing Bob's message. In the positive to more positive condition, Frank reported he was originally "somewhere around a $+1$ " but shifted to "somewhere around a $+9$." All participants were explicitly informed that "Frank's position moved about 8 points up" on the scale following his interaction with Bob.

8.1.4. Perceived change

Our primary measure was the same index of perceived change as in Experiments 1–3 ($\alpha = 0.94$). Participants also completed an exploratory personality index (see Online supplement).

8.2. Results and discussion

We found a significant effect of change condition on perceived change, $F(2, 299) = 4.62$, $p = .01$, $\eta^2 = 0.03$ ($d = 0.35$; see Fig. 1). As predicted, change from negative to positive ($M = 5.56$; $SD = 1.79$) was perceived as greater than change from negative to less negative ($M = 4.78$; $SD = 1.97$), $t(299) = -2.83$, $p = .005$, $d = -0.41$, and change from positive to more positive ($M = 4.91$; $SD = 2.00$), $t(299) = -2.39$, $p = .02$, $d = -0.34$. Unlike the earlier results, these latter two conditions did not differ from each other, $t(299) = -0.47$, $p = .64$, $d = -0.06$. Most germane, despite several methodological changes, Experiment 4 provided further support for the qualitative change hypothesis.

9. Experiment 5

Experiment 5 had multiple objectives. First, we assessed the

mechanism driving the qualitative change effect. As outlined earlier, we posit that this effect is driven by processing ease. More specifically, we submit that people find it easier to perceive and understand attitude change when it occurs across valence rather than within valence. This ease, in turn, is predicted to amplify perceptions of the amount of change that has occurred. To explore this mechanism in Experiment 5, we measured ease of processing using items adapted from past research (e.g., Labroo & Lee, 2006; Lee & Aaker, 2004; Schwarz, 2004).

In addition, we examined an alternative account for the qualitative change effect. It could be that instead of (or in addition to) ease of processing, the effect is driven by *extent* of processing. Perhaps qualitative changes capture more attention and elicit greater interest, involvement, and thought than do non-qualitative changes. If true, qualitative changes might seem greater as a function of how much attention and thought people direct toward them. Indeed, ample evidence from the attitudes literature suggests that increased thinking about something can amplify people's reactions to that thing. For example, research on the mere thought effect reveals that just thinking about an attitude object can cause attitude polarization, because people tend to think attitude-consistent thoughts, which push their attitudes to become more extreme (e.g., Chaiken & Yates, 1985; Clarkson et al., 2011; Tesser, 1978). Similarly, in the persuasion literature, greater thinking about a persuasive message can enhance that message's impact (e.g., Karmarkar & Tormala, 2010; Petty & Cacioppo, 1986). Analogously, perhaps when people note that attitude change has occurred, more thought about that change (which might occur in the qualitative change context) augments perceptions of how much changed has occurred. To assess this possibility, we measured self-reported processing in Experiment 5.

Another goal of Experiment 5 was to test a consequence of the qualitative change effect. Given attitudes researchers' enduring interest in attitude-behavior correspondence over the years (e.g., Fazio & Olson, 2014; Glasman & Albarracín, 2006; Wicker, 1969), we examined the consequences of perceived attitude change for assessments of the behavioral implications of that change. In essence, we were interested in the possibility that because people see qualitative change as larger in magnitude than non-qualitative change, they might also see it as more meaningful or significant; for example, as having more impact on behavior. To test this possibility, we measured inferences about the behavioral implications of the attitude change observed.

Finally, to further establish the robustness of the qualitative change effect, we made several modifications to the experiment. Specifically, we recruited participants from a different online panel and presented them with a new scenario that did not involve Bob, Frank, and an abstract policy. Instead, the scenario described attitude change in a self-persuasion context with a concrete issue that participants themselves might be thinking about. Specifically, participants read about an individual named Ashley who changed her attitude after thinking about the benefits and drawbacks of self-driving cars. We hypothesized that participants would perceive qualitative change as significantly greater than non-qualitative change, and that perceptions of attitude change would influence predictions of subsequent behavior change.

9.1. Method

9.1.1. Participants and design

A total of 616 participants from the online Qualtrics panel took part for monetary compensation.³ Participants were randomly assigned to

³ Participants were recruited over two virtually identical waves (1st wave $n = 317$; 2nd wave $n = 299$). The only difference was that in the first wave we included an additional measure assessing the ease with which participants understood the scenario and asked participants to report their thoughts about the scenario (see Online supplement for measures and results). The purpose of the second wave was to have sufficient power to adequately test our

one of three change conditions: negative to less negative, negative to positive, or positive to more positive. A sensitivity analysis revealed 80% power to detect a minimum effect size of $\eta^2 = 0.02$ ($d = 0.29$). This sample also was sufficient for testing the mediating role of processing ease on perceived change with 80% power (Fritz & MacKinnon, 2007).

9.1.2. Procedure

Participants read about an individual, Ashley, who had been thinking about the benefits and drawbacks of self-driving cars. Ashley had an initial opinion of self-driving cars, but over a period of a few months she had been thinking about the arguments in favor of and against them, and she shifted her opinion. Participants received information about Ashley's initial and final attitude on a scale ranging from -50 to $+50$, where -50 meant "strongly opposed," $+50$ meant "strongly in favor," and 0 meant "neutral." Following the scenario and perceived change index, participants completed a series of measures assessing processing ease, extent of processing, and behavioral implications.

9.1.3. Change condition

In the negative to positive condition, participants read that Ashley was at first "somewhat opposed" to self-driving cars, but that she shifted and had become "somewhat in favor" after thinking more, moving from -6 to $+6$ on the attitude scale. In the negative to less negative condition, participants read that Ashley was "pretty opposed" but shifted to become "less opposed," moving from -18 to -6 on the scale. In the positive to more positive condition, participants read that Ashley was "somewhat in favor" but shifted to become "more in favor," moving from $+6$ to $+18$ on the scale. In each condition, we explicitly noted that "Ashley's position moved 12 points up" on the attitude scale.

9.1.4. Dependent measures

9.1.4.1. Perceived change. Our primary measure was an index of four items assessing perceived change. Because Ashley was the lone character in the scenario, all four items assessed perceptions of her change. On scales ranging from 1 to 9, participants rated how much Ashley changed her opinion, how much Ashley shifted her view over time, how big a change Ashley's thoughts caused in her stance on the issue, and how different Ashley's new view was compared to her old view. These four measures were averaged ($\alpha = 0.93$); higher values reflected perceptions of greater change.

9.1.4.2. Processing ease. The measure of processing ease included two items. On scales ranging from 1 to 9, participants rated how easy it was to understand the way Ashley's attitude changed and how easy it was to see the change in Ashley's view before and after she thought about the issue. Past research has successfully measured processing ease using subjective measures (Schwarz, 2004), and these specific measures were adapted from past research (e.g., Labroo & Lee, 2006; Lee & Aaker, 2004; Mayer & Tormala, 2010). Responses were averaged ($r [614] = 0.80$, $p < .001$); higher values reflected greater processing ease.

9.1.4.3. Extent of processing. Our measure of extent of processing included three items assessing involvement, interest, and depth of thinking. The involvement measure was adapted from Lee and Labroo (2004), and asked participants to report how involved they were while reading the scenario (ranging from "skimmed it quickly, not at all involved" to "paid a lot of attention, very involved"). For the interest

(footnote continued)

hypothesized mechanism and possible alternative explanations (Fritz & MacKinnon, 2007, p. 237). The procedural changes to the second wave were made to refine our measure of ease and reduce the burden on participants.

and depth of thinking measures, participants rated how interesting the scenario was and how deeply they thought about the information in the scenario, respectively. These measures were also adapted from past research (e.g., Karmarkar & Tormala, 2010; Petty, Briñol, & Tormala, 2002; Tormala, Petty, & Briñol, 2002). All ratings were made on scales ranging from 1 to 9, with higher values reflecting greater involvement, interest, and depth of thinking. Responses were averaged ($\alpha = 0.85$).

9.1.4.4. Behavioral implications. Finally, we included two items assessing the extent to which participants viewed Ashley's attitude change as having implications for her future behavior. On scales ranging from 1 to 9, participants indicated how likely it was that the change in Ashley's view would affect her future behavior (e.g., whether she would purchase a self-driving car or whether she would ride in a self-driving car) and the extent to which the change in Ashley's view meant her behavior would change as well. Responses were averaged ($r [614] = 0.78, p < .001$); higher values reflected greater predicted behavioral implications.

9.2. Results

We submitted each index to a one-way ANOVA with change condition as the independent variable. When this analysis revealed a statistically significant effect of change condition, we conducted planned comparisons. Means and standard deviations are in Table 2.

9.2.1. Perceived change

There was a significant effect of change condition on perceived change, $F(2, 613) = 15.41, p < .001, \eta^2 = 0.05 (d = 0.46)$, such that change from negative to positive was perceived as greater than change from negative to less negative, $t(613) = -5.50, p < .001, d = -0.56$, and change from positive to more positive, $t(613) = -3.41, p < .001, d = -0.33$. Participants also perceived that Ashley changed more when she moved from positive to more positive rather than from negative to less negative, $t(613) = -2.07, p = .04, d = -0.20$.

9.2.2. Processing ease

Change condition also affected processing ease, $F(2, 613) = 4.64, p = .01, \eta^2 = 0.01 (d = 0.20)$. Change from negative to positive was easier to process than change from negative to less negative, $t(613) = -2.71, p = .01, d = -0.28$, and change from positive to more positive, $t(613) = -2.57, p = .01, d = -0.25$, which did not differ from each other, $t(613) = -0.13, p = .90, d = -0.01$.

Table 2

Dependent variable measures in Experiment 5. Subscripts should be interpreted within rows only. Values with the same subscript do not differ from each other at a $p < .05$ level. † indicates values with the same subscript differ at a $p < .10$ level.

Dependent measures	Change condition		
	Negative to less negative	Negative to positive	Positive to more positive
Perceived change			
M	5.24 _a	6.16 _b	5.59 _c
SD	1.61	1.69	1.81
Processing ease			
M	6.40 _a	6.89 _b	6.42 _a
SD	1.81	1.73	1.98
Extent of processing			
M	6.13 _a	6.37 _a	6.17 _a
SD	1.92	1.82	1.93
Behavioral implications			
M	5.70 _a	6.42 _b	6.03 _{a†}
SD	1.98	1.71	1.93

9.2.3. Extent of processing

There was no effect of change condition on self-reported extent of processing, $F(2, 613) = 1.00, p = .37, \eta^2 = 0.003 (d = 0.11)$.

9.2.4. Behavioral implications

There was a significant effect on behavioral implications, $F(2, 613) = 7.55, p < .001, \eta^2 = 0.02 (d = 0.29)$. Change from negative to positive was perceived as implying greater future behavior change than was change from negative to less negative, $t(613) = -3.88, p < .001, d = -0.39$, or change from positive to more positive, $t(613) = -2.09, p = .04, d = -0.21$. Participants also perceived that change from positive to more positive was marginally more predictive of future behavior change than change from negative to less negative, $t(613) = -1.79, p = .07, d = -0.17$.

9.2.5. Mediation

Our initial analyses suggested that processing ease was a possible mediator of the qualitative change effect, but that extent of processing was not. We also predicted that augmented perceptions of attitude change would increase predictions of subsequent behavior change. To test this prediction and determine whether processing ease mediated the effect of change condition on perceived change, we conducted a bootstrapped serial mediation analysis (Hayes & Preacher, 2014) that tested the pathway from condition to processing ease to perceived change to behavioral implications. Because our independent variable was multicategorical (three levels), we created two orthogonal contrasts (Hayes & Preacher, 2014). The first contrast compared the qualitative change condition to the two non-qualitative change conditions; the second compared the two non-qualitative change conditions to each other. We tested for mediation of the first contrast controlling for the second, and of the second contrast controlling for the first.

Most pertinent to our primary concerns was assessing the mechanism for the qualitative change effect—that is, testing whether processing ease mediated the first contrast (controlling for the second). As a secondary interest, in the same mediation model we gauged whether the qualitative change effect had implications for perceptions of future behavior change. We found support for this mediation prediction. Processing ease mediated the effect of condition on perceived change, which in turn influenced participants' predictions of future behavior change (95% CI = [0.013, 0.063]; Fig. 2). Although not germane to our primary concerns, a mediation analysis for the comparison between the two non-qualitative change conditions (i.e., the second contrast, controlling for the first) was not significant (95% CI = [-0.043, 0.041]).

9.3. Discussion

Experiment 5 replicated the key results from Experiments 1–4 in a new context. Again, qualitative change was perceived as greater than non-qualitative change. We also extended these findings in two major ways. First, we showed that processing ease mediated the effect. Participants found it easier to process qualitative changes than non-

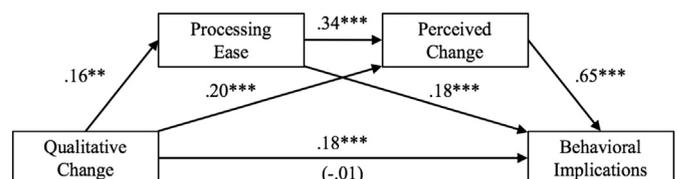


Fig. 2. Mediation Analysis in Experiment 5 for the comparison of the qualitative change condition to the two non-qualitative change conditions (controlling for the comparison between the two non-qualitative change conditions). The path coefficients are unstandardized betas. The value in parenthesis indicates the effect of change condition on behavioral implications when controlling for the mediators. * $p < .05$ ** $p < .01$ *** $p < .001$.

qualitative changes, and this ease predicted perceptions of change. Second, we highlighted a consequence of the qualitative change effect: Predicted changes in subsequent behavior were greater following qualitative rather than non-qualitative change. Together, these effects produced serial mediation: Qualitative change was easier to process, which increased its perceived magnitude and, in turn, fostered increased perceptions of future behavior change. This finding suggests that people see qualitative change as both larger and more meaningful, or significant, than non-qualitative change.

Of note, we also examined alternative mediation models: a reverse serial model in which perceived behavioral implications mediated the perceived change effect (qualitative change → ease → behavioral implications → perceived change), a reverse serial model in which perceived change mediated the ease and behavioral implications effects (qualitative change → perceived change → ease → behavioral implications), and a parallel model in which ease and behavioral implications simultaneously mediated the effect of qualitative change on perceived change. Each model produced evidence of successful mediation. In the first, we found a significant indirect effect from condition (first contrast, controlling for the second) to ease to behavioral implications to perceived change (95% CI = [0.012, 0.062]). In the second, there was a significant indirect effect from condition (first contrast, controlling for the second) to perceived change to ease to behavioral implications (95% CI = [0.009, 0.031]). In the third, the parallel model, condition (first contrast, controlling for the second) simultaneously affected perceived attitude change through both ease (95% CI = [0.005, 0.040]) and behavioral implications (95% CI = [0.047, 0.161]).

Despite their general statistical equivalence, there is reason to favor the hypothesized model over these alternatives. First, as reviewed earlier, it is grounded in past research on categorical perception and processing fluency. Logically too, the hypothesized model holds up better than the alternative models (see Pieters, 2017). Indeed, the first and third models suggest that participants perceived the target to have changed her attitude more because her attitude change meant her behavior would change to a greater extent in the future. We find this pathway less plausible than the hypothesized one based on the manipulation and measures employed. Again, participants were exposed to information describing the target's attitude change and subsequently rated her perceived change and indicated the extent to which her attitude change would lead to future behavior change. It seems unlikely that participants read about her attitude change, inferred that this attitude change would cause her future behavior to change, and only then concluded that her attitude must have changed a great deal. It is more plausible that participants read about the target's attitude change, judged that change to be large or small, and then drew inferences about what this meant for her future behavior. To be sure, perceived change and inferred behavioral implications were highly correlated and likely influence each other in both directions, but we view the proposed mediating sequence as more plausible than the reverse in this instance.

The second alternative model also has logical limitations. That model suggests that perceived attitude change causes processing ease, which leads to inferences of greater subsequent behavior change. This seems psychologically inefficient in requiring that people perceived a given amount of change, then decided how easy that change must have been to process, and only then drew inferences about behavioral implications. Processing ease and perceived change were correlated and may indeed influence each other in both directions, but we see this pathway as psychologically inefficient, and as leaving open the question of what fosters greater (or less) perceived change in the first place. Ultimately, although the data from Experiment 5 do not clearly favor one mediation model over the others, we favor the hypothesized model on conceptual grounds.

As a final note, we found no evidence for the notion that *extent* of processing played a role in the qualitative change effect. Is it possible that our measures of extent of processing were imperfect and failed to fully capture participants' processing levels? Although we cannot rule

this out based on Experiment 5 alone, it is worth noting that we adopted our measures from past research, where they have been used successfully. Indeed, asking participants to self-report the extent to which they have processed information is a common and recommended procedure in attitudes and persuasion research (e.g., Petty et al., 2002; Petty & Cacioppo, 1979). Moreover, prior experiments have shown that perceived and actual processing are positively correlated, and that perceived processing can mediate the relationship between actual processing and other outcomes (e.g., Barden & Petty, 2008; Barden & Tormala, 2014). Thus, we see it as unlikely that the null effect for extent of processing stems from issues related to our self-report measures.

10. Experiment 6

Experiment 6 had several goals. First, we sought to test a concrete consequence of the qualitative change effect. Specifically, we examined whether perceived change could affect interpersonal perceptions. To do so, we moved the experiment into a political context. We reasoned that the qualitative change effect could have implications for how people think about politicians who change their views during their campaigns. Past research suggests that people dislike it when politicians change their opinions over time, particularly when doing so is viewed as strategic (i.e., to capture votes) rather than principled (e.g., Hoffman, Yoeli, & Nowak, 2015; see also Kreps et al., 2017; Reich & Tormala, 2013). Voters' dislike for politicians who change their views is thought to stem at least partly from a general dislike of inconsistency (Tomz & Van Houweling, 2009). Based on these findings, and the results of Experiments 1–5, we predicted that participants would express a reduced likelihood of voting for politicians whose opinions changed qualitatively rather than non-qualitatively over the course of their campaigns.

As a second objective, we sought to test the qualitative change effect in a joint (within-participant) evaluation task. That is, participants evaluated multiple targets simultaneously, each of whom changed an identical degree, but we varied the location of the targets' change on an attitude scale. Specifically, participants read about three politicians whose opinions had changed during their campaigns. One changed qualitatively (from positive to negative in this experiment) whereas the others changed non-qualitatively (from positive to less positive or negative to more negative). We asked participants which politician they felt changed the most and which they would be least likely to vote for. This paradigm offered a conservative test of our hypothesis as participants could plainly see that each politician had changed the same amount on an attitude scale. To strengthen the test, we allowed participants to choose any of the politicians or indicate that all three changed the same amount. If we replicated our findings in this context, it would suggest that the qualitative change effect supersedes the objective observation of identical absolute change across targets.

As a final modification to establish the robustness of the qualitative change effect, we altered the attitude scale along which change was described. In this experiment, we used a 9-point attitude scale, which is an extremely common one in the attitudes literature (e.g., Albarracín & Handley, 2011; Blankenship et al., 2012; Clark & Evans, 2014; Clarkson et al., 2011; Rucker & Petty, 2004; Tormala & Petty, 2002). Using a 9-point scale allowed us to test the qualitative change effect in a widely-employed scale format.

10.1. Method

10.1.1. Participants and procedure

A total of 152 participants from Amazon's Mechanical Turk took part for monetary compensation. A sensitivity analysis revealed 80% power to detect an effect size of $w = 0.23$.

All participants read about three politicians running for office. Participants were informed that each politician had stated a policy position on a scale ranging from 1 (strongly opposed) to 9 (strongly in favor) at the beginning of the campaign. However, over the course of

the campaign, each politician changed his or her attitude toward the policy and became more opposed to it, shifting two points down on the scale. Specifically, participants learned that Politician #1 was “somewhat opposed” to the policy at the beginning of the campaign and even “more opposed” at the end, Politician #2 was “somewhat in favor” at the beginning and “somewhat opposed” at the end, and Politician #3 was “pretty in favor” at the beginning and “less in favor” at the end. These changes corresponded to shifts from 4 to 2, 6 to 4, and 8 to 6 on the attitude scale, respectively. Following the scenario, participants completed dependent measures and an attention check.

10.1.2. Dependent measures

10.1.2.1. Perceived change and consistency. Immediately after the scenario, participants answered two questions about which politician changed the most: “In your personal opinion, which candidate do you feel changed his or her opinion the most?” and, “Which politician had the least consistent view over time?” Participants responded by selecting among the three politicians, or by indicating that all three politicians changed equally or were equally consistent. Participants also answered an exploratory question about the politicians (see Online supplement).

10.1.2.2. Voting likelihood. Participants then indicated which politician they would be *least* likely to vote for. Again, participants were not forced to choose among the politicians; they could indicate that they would be equally likely to vote (or not vote) for each politician.

10.1.3. Attention check

Because the voting likelihood measure was negatively framed, asking which politician participants would be *least* likely to vote for, we included an attention check to make sure participants understood the question and did not incorrectly read it as asking who they would be *most* likely to vote for. Specifically, at the end of the experiment, participants were shown two questions and asked to select which of the questions they had been asked previously. The options were “Which politician would you be least likely to vote for?” (the correct answer) and “Which politician would you be most likely to vote for?” We made an a priori decision to exclude any participants who answered this question incorrectly.

10.2. Results

One hundred fourteen participants (75%) passed the attention check, correctly selecting the question they were previously asked. The analyses presented below exclude participants who failed the attention check. The significance of the results does not change when these participants are included.

As illustrated in Fig. 3, on each measure participants were more likely to select the politician who had changed qualitatively than either of the other politicians. To determine if this effect was significant, we excluded participants who chose the “all the same” option and conducted a chi-square goodness-of-fit test in which the expected probabilities of selecting the politician who changed qualitatively (Politician #2) or a politician who did not change qualitatively (Politician #1 or Politician #3) were equal. That is, the null was that there was a 50% chance of choosing Politician #2 and a 50% chance of choosing Politician #1 or Politician #3, given that a specific politician was selected. This analysis offered a conservative test of the qualitative change effect, because choice share for Politician #2 (who changed qualitatively) would have to exceed 50% (i.e., 50% of participants who did not select “all the same”) to obtain a significant effect. We made no predictions about the likelihood of choosing Politician #1 versus Politician #3. Alternative chi-square tests comparing choice frequencies of Politician #1, Politician #2, and Politician #3 to expected probabilities of 33% each (after excluding “all the same” responses) are also significant for each of the results reported below.

10.2.1. Perceived change and consistency

For the question of which politician changed most, results indicated that the choice distribution differed from chance, $\chi^2(1, n = 114) = 31.84, p < .001, w = 0.58$. Among those who chose a politician, 79% of participants saw the politician who changed qualitatively as having changed the most (see Fig. 3 for choice distribution for the full sample including those who selected “all the same”). We found a similar pattern for the perceived consistency measure: 76% of participants who chose a politician indicated that the politician who changed qualitatively was the least consistent over time, $\chi^2(1, n = 114) = 25.51, p < .001, w = 0.51$.

10.2.2. Voting likelihood

We submitted voting responses to the same chi-square goodness-of-fit test. The observed frequencies significantly differed from chance, $\chi^2(1, n = 114) = 14.10, p < .001, w = 0.41$, indicating that 71% participants who chose a politician would be least likely to vote for the politician who changed qualitatively. To gauge the extent to which perceived change predicted the voting outcome, we calculated Fleiss's Kappa (Fleiss, 1971) to determine the level of consistency across the perceived change, consistency, and voting measures. Results revealed a Kappa statistic of 0.51 across the three measures, suggesting that viewing a politician as having changed the most and being the least consistent over time predicted being least likely to vote for that politician with moderate strength (see Landis & Koch, 1977).

10.3. Discussion

Experiment 6 again revealed that a qualitative shift accentuated perceived change. This change was also consequential; participants' reported likelihood of voting for a politician reduced when that politician's attitude changed qualitatively rather than non-qualitatively. This effect emerged despite the fact that all three targets changed an identical amount in a within-participant design. That the effect emerged in this context speaks to the power of qualitative shifts in increasing perceived change even when absolutely equivalent non-qualitative shifts are plainly apparent. As shown in Fig. 3, only a small minority of respondents selected the “all the same” option.

One concern that could be raised is that participants' responses to the voting question might have been influenced by their responses to the perceived change and consistency measures. To assess this possibility, we conducted a follow-up in-lab study in which 110 students and staff at a west coast university received the same scenario and voting measure (and attention check; $N_{\text{passing}} = 86$; 78% pass rate) without the perceived change or consistency questions. This experiment replicated the vote distribution results of Experiment 6, $\chi^2(1, n = 86) = 5.23, p = .02, w = 0.28$, such that 64% of respondents who chose a politician (and 51% of the full sample, including those who selected the “all the same” option) reported being least likely to vote for the politician who changed qualitatively.

Finally, one might wonder whether including an “all the same” option ironically boosted our effect by highlighting the fact that change is objectively equivalent and perhaps calling more attention to valence than normally would exist. In two additional experiments (total $N = 599$), we replicated the basic design of Experiment 6 but changed aspects of the context and asked the question, “Which politician had the least consistent view over time?” while removing the “all the same” option. Across these experiments, we found that 69% of all participants perceived that the politician who changed qualitatively was the least consistent over time (χ^2 s $> 37.96, ps < .001, ws > 0.36$). Thus, including the “all the same” option in Experiment 6 does not appear to boost the choice share for the politician who changed qualitatively among the full sample of participants.

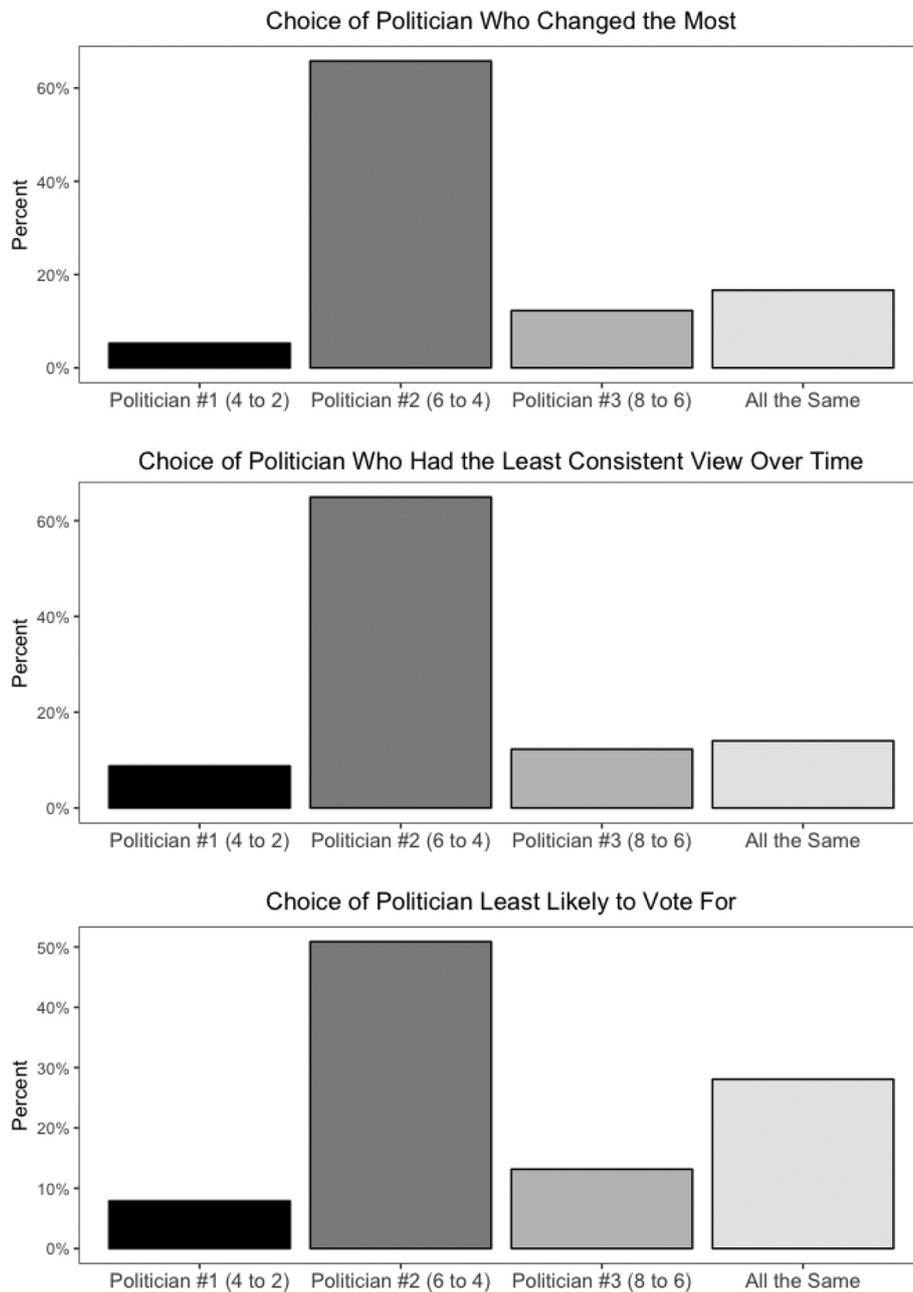


Fig. 3. Percent of participants choosing each politician for the measures in Experiment 6. Top panel: In your personal opinion, which candidate do you feel changed his or her opinion the most? Middle panel: Which politician had the least consistent view over time? Bottom panel: Which politician would you be least likely to vote for?

11. General discussion

Despite a vast literature yielding countless insights into attitude change as a psychological phenomenon, as a field we know very little about the factors that shape people's *perceptions* of attitude change—in particular, the variables that influence whether we see an individual's attitude change as large or small. The current research represents an initial foray in this domain. Although several distinct hypotheses were possible, we drew upon past research on categorical perception and processing fluency and proposed that people would perceive qualitative change—in this case, change in attitude valence—as greater in magnitude than non-qualitative change—that is, change *within* attitude valence.

Across multiple experiments we found consistent support for this hypothesis. First, in Experiment 1, we observed the effect in the absence

of an attitude scale, which suggests that the effect is not merely an artifact of people's perceptions of ordinal scales and the distance between individual scale points. In addition, we observed the effect across variations in the *type* of attitude scale presented (e.g., 0 to 100, –50 to 50, or 1 to 9), the direction of change described, the absolute amount of change described, the relative location of the change along an evaluative continuum, the attitude change scenario or context, and the abstractness or concreteness of the attitude issue. In each case, we found that qualitative attitude change was perceived as greater than non-qualitative attitude change. This effect was mediated by ease of processing and had consequences for other important outcomes. Specifically, participants predicted greater behavior change when they read about qualitative rather than non-qualitative attitude change and they indicated that they were less likely to vote for a politician whose attitude changed qualitatively compared to non-qualitatively.

Are people correct in perceiving qualitative change as greater, and more impactful, than non-qualitative change? This question goes beyond the scope of the current research, which sought to document the effect of qualitative versus non-qualitative shifts on perceived attitude change, but we note that there are two salient interpretations of our findings. One is that they reflect reality—that is, that changing valence actually *is* greater and more impactful than changing within valence. After all, if one's attitude changes valence, it would seem to guide behavior in a whole new direction. However, it could be that people are incorrect—that is, that they are making an error—and that in reality changes in attitude extremity tend to be more impactful. As outlined earlier, an attitude strength perspective suggests that moving people from somewhat positive to extremely positive, for example, could alter their behavior more than moving them from somewhat negative to somewhat positive. Indeed, in the latter case people might not act on their attitude either way as they are only leaning slightly. Thus, based on the attitude strength literature, changes in attitude valence might actually have less impact on behavior than changes within attitude valence. Our goal in this initial work was not to establish that people are or are not making an error in seeing qualitative change as greater than non-qualitative change; we simply aimed to test the qualitative change hypotheses against two plausible rivals (the null and extremity hypotheses). However, investigating this issue systematically in future work would be worthwhile.

Error or not, the qualitative change effect has potentially important implications. First, at a broad theoretical level, our studies ask and answer a novel question: What influences people's perceptions that attitude change is large or small in degree? We identify one factor—whether an evaluative shift in valence occurs—that strikes at the heart of attitudes and persuasion research. We see this finding as expanding our insight into how people evaluate change in a wide variety of contexts. Imagine an ad campaign showcasing how much a company has improved its product (e.g., Domino's "Pizza Turnaround" campaign). Our findings suggest that consumers might be most impressed by the degree of change described when it involves a change of valence (from bad to good) rather than a change within valence (from moderate to very good). As another example, imagine a customer service strategy that enhances customer satisfaction with a particular airline. Which type of enhancement would the airline's executives view as the greatest? Our findings suggest that decision makers might see a shift from dissatisfied to satisfied as bigger than a shift from satisfied to very satisfied. Across diverse contexts (e.g., improved ratings for a wine from one vintage to the next, improved teaching evaluations from one academic year to the next), our findings point to potential differences in change magnitude estimates depending on where change occurs along an evaluative continuum.

The current research also has interpersonal implications. As one example, the qualitative change effect could influence how people evaluate others who change their own attitudes or attitude-relevant behaviors over time. Consider the role of change in fostering inspiration (Klein & O'Brien, 2017). Could it be that individuals find others more inspiring when those others change from bad to good rather than from bad to less bad or good to more good? Our research suggests so. The current studies also highlight potential implications for how people react to hypocrisy or flip-flopping in others. As demonstrated in Experiment 6, people may punish others more when they express views that differ in valence as opposed to differing within valence.

Finally, our studies have methodological implications. As discussed earlier, the dominant method for assessing attitude change in the literature is to measure attitudes before and after some treatment or manipulation and compute the difference. This approach is sensible, but our findings suggest that researchers' calculations of attitude change might not map onto the way people perceive and think about change in their daily lives. In short, if people perceive change to be greater when it involves a shift in valence, perhaps researchers should refine their calculations of attitude change to take this differential

weighting into account. Doing so could reduce the disconnect between how researchers interpret attitude change and how laypeople perceive it.

11.1. Remaining questions and future directions

The current research provides initial insight into an understudied area of attitudes research. At the same time, it raises new questions, which we hope will stimulate other researchers to help expand our understanding of perceived attitude change. Here, we summarize what we see as some of the next steps in this area.

11.1.1. Valence asymmetry

First, although it was not the focus of our research, we uncovered a valence asymmetry in our experiments: Attitude change in the positive domain was generally perceived as greater than attitude change in the negative domain, regardless of whether that change moved toward or away from neutrality. This effect was not predicted, but it is interesting and reliable and it could prove important to our understanding of attitude change perceptions. At this point, we can only speculate as to why this asymmetry emerged, though Experiment 5 suggests that it does not stem from the ease or extent of processing. One possibility is that due to negativity bias—the disproportionate weight given to negative over positive information (Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001; Cacioppo, Gardner, & Berntson, 1997; Rozin & Royzman, 2001)—negative attitudes are construed as *essentially* negative regardless of whether they recently became more or less negative. That is, perhaps both moderate and extreme negative attitudes carry psychological weight, reducing the perceived difference between Time 1 and Time 2 negative attitudes, whereas only extreme positive attitudes carry psychological weight, *accentuating* the difference between Time 1 and Time 2 positive attitudes. This account is purely speculative, but it might offer one explanation for valence asymmetry in perceived attitude change. For now, we note that this finding underscores the notion that attitude change is viewed differently depending on where along an evaluative continuum it occurs, and we call for further investigation on valence asymmetries in perceived attitude change.

11.1.2. Perceiving one's own change

Another promising direction for future research would be to explore whether people experience *their own* attitude change as greater when it is qualitative rather than non-qualitative in nature. Each of our experiments examined perceptions of *others'* attitude change. We focused on others' change because doing so enabled us to design experiments that carefully controlled the amount and direction of attitude change across experimental conditions. Nevertheless, the question of how people perceive their own attitude change is an interesting and important one.

To offer a first look, we conducted a follow-up study. In this study, 199 participants (recruited using the Lucid Fulcrum panel) reported their attitudes toward self-driving cars on a scale ranging from -50 to 50 , where -50 meant "against," 50 meant "in favor," and 0 meant "neutral." After reporting their initial attitudes, participants read a passage containing counterattitudinal arguments about self-driving cars. Following this message, participants indicated their current attitudes toward self-driving cars on the same scale as before and then completed a measure of perceived change: "How much did reading those arguments change your attitude on self-driving cars?" (1 [not at all] to 9 [very much]).

We used participants' Time 1 and Time 2 attitudes to categorize them into two groups: qualitative change ($n = 81$) and non-qualitative change ($n = 82$). This determination was based on whether participants' Time 1 and Time 2 attitudes spanned the neutral scale midpoint. Due to the difficulty in classifying participants who displayed zero change or shifted to or from the neutral midpoint of the attitude scale, we made an a priori decision to exclude them ($n = 36$). Then, to assess

the relationship between change type (qualitative versus non-qualitative) and perceived change, we regressed perceived attitude change onto change type controlling for objective change magnitude (i.e., the absolute difference between Time 1 and Time 2 attitudes). As expected, objective change magnitude predicted perceived change ($\beta = 0.04$, $SE = 0.01$, $t(160) = 3.77$, $p < .001$), such that greater absolute change predicted greater perceived change. Most germane to our primary concerns, change type also predicted perceived change ($\beta = 1.21$, $SE = 0.50$, $t(160) = 2.44$, $p = .02$); perceived change was greater when the change was qualitative ($M_{\text{adjusted}} = 6.74$) rather than non-qualitative ($M_{\text{adjusted}} = 5.53$).⁴ There was no interaction between change type and initial attitude ($\beta = 0.82$, $SE = 0.62$, $t(158) = 1.33$, $p = .19$).

These results suggest that in addition to perceiving others' attitude change as greater when it is qualitative in nature, people experience their own attitude change as greater when their attitudes shift qualitatively, controlling for the absolute magnitude of their change. Also noteworthy, this study demonstrates the qualitative change effect in a paradigm in which participants received a message that actually changed their attitudes. One potential limitation of Experiments 1–6 is that although they uncovered the predicted effect despite numerous procedural modifications, each study employed a scenario paradigm in which we explicitly described the type and often the amount of change that occurred. Although people in the wild often do receive explicit information about individuals' attitudes and attitude change, they frequently must observe or infer this change more spontaneously. Thus, this final study pushes our findings into a context in which change is not described explicitly but rather observed—or *experienced*—spontaneously. That the effect emerges in this context speaks to the robustness of the phenomenon and helps establish its ecological validity.

In future research, it would be useful to follow up on this finding to determine how perceiving one's own change affects one's future behavior. One straightforward hypothesis would be that people align their behavior change with their perceptions of attitude change, such that the more people believe their attitudes have changed, the more they change their behavior as well. On the other hand, maybe the perception of extensive attitude change triggers uncertainty, which hinders behavior change (see Tormala & Rucker, 2018). Perhaps when people think they have changed their attitude a great deal, for instance, they proceed with caution before engaging in new or different behavior, waiting to re-evaluate the issue or their reaction to it. If true, perceived attitude change negatively correlates with actual behavior change. We see this as an intriguing issue worth exploring in future work.

11.1.3. Defining persuasion success

The current research also raises questions pertaining to how people define persuasion, and what they believe constitutes a successful persuasion outcome. Based on our findings, it is reasonable to wonder if people view the goal of persuasion to be changing the valence of people's attitudes—as moving them from one side of an issue to the other. If so, non-qualitative attitude change would not constitute successful persuasion, by definition, and only the qualitative change conditions would have met the criteria for successful persuasion in the current studies. This issue warrants scrutiny.

To be sure, it could indeed be that when people think about persuasion, the goal they bring to mind involves qualitative change. This response could be an instantiation of the qualitative change effect itself. However, we do not believe the qualitative change effect is contingent upon this *definition* of persuasion or attitude change. As one piece of evidence, participants in the current experiments generally perceived moderate to high degrees of change (i.e., change scores at or above the midpoint of the unipolar scales), even in the non-qualitative change

conditions. As an exploratory measure, Experiments 1–3 also assessed participants' perceptions of how successful the source was in changing the recipient's attitude (see Online supplement; items 1.2, 2.2, and 3.2). While these results largely parallel the perceived change findings, participants viewed the source as moderately to highly successful across conditions, even when the change described was non-qualitative. Thus, qualitative change was not necessary for participants to see the attitude change effort as least somewhat successful.

Nonetheless, the question of how people define successful persuasion, and how this definition affects perceived change, is an interesting one worth exploring in future research. Ultimately, we suspect that both the goal and definition of persuasion varies by context. For instance, whereas the typical goal of persuasion might be to shift valence in a product or policy context, it might differ in a health behavior or environmental action context, where influencers might seek to elicit *more* of a particular behavior rather than a qualitatively different behavior. Future work would profit from further consideration of these issues.

11.1.4. Other types of qualitative change

Finally, although we focused on understanding the effect of changing valence versus changing within valence on the perceived magnitude of attitude change, we believe other forms of qualitative versus non-qualitative change are worth investigating in future research. As one example, perhaps people perceive greater change when someone's attitude changes in basis—for example, when it shifts from an affective to a cognitive attitude (see Fabrigar & Petty, 1999) or from a moral to a non-moral attitude (see Luttrell, Petty, Briñol, & Wagner, 2016)—compared to when that attitude shifts *within* basis (e.g., becoming even more emotional or even more moralized). Even when the qualitative threshold is arbitrary, or artificial, we would expect to observe this sort of effect. For instance, if a professor needs to receive an average teaching rating of 4 (on a scale ranging from 1 to 5) to have her appointment renewed, we might expect a change from 3.5 to 4.5 to be perceived as greater, and more meaningful or impactful, than change from 2.5 to 3.5, even if 3 represented a neutral midpoint on a bipolar evaluation scale. Attitude valence and extremity are core dimensions of the attitude construct, so we started here in this initial research, but other forms of qualitative change are worth studying in future research.

11.2. Conclusion

The current research offers an initial look at a new question: What shapes people's perceptions that attitude change is large or small? From a theoretical and practical perspective, this question speaks to a fundamental aspect of attitude change. How do we perceive it? What affects its perceived magnitude? How do we react to change in our own and others' views? As an initial foothold, we investigated the role of qualitative versus non-qualitative change. We found a robust effect whereby people perceive greater change when that change is qualitative in nature. However, we see this as a first step in a new direction for attitudes research. Much remains much to be learned and we hope other researchers will join our call to better understand the psychology of perceived attitude change.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jesp.2019.02.001>.

References

- Abelson, R. P. (1995). Attitude extremity. In R. E. Petty, & J. A. Krosnick (Eds.). *Attitude strength: Antecedents and consequences* (pp. 25–41). Mahwah, NJ: Erlbaum.
- Abelson, R. P., Kinder, D. R., Peters, M. D., & Fiske, S. T. (1982). Affective and semantic components in political person perception. *Journal of Personality and Social*

⁴ Adjusted means control for actual change and are at the mean level of actual change.

- Psychology*, 42(4), 619.
- 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.
- Albarracín, D., & Handley, I. M. (2011). The time for doing is not the time for change: Effects of general action and inaction goals on attitude retrieval and attitude change. *Journal of Personality and Social Psychology*, 100(6), 983.
- Alter, A. L., & Oppenheimer, D. M. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychology Review*, 13(3), 219–235.
- Attneave, F. (1955). Symmetry, information, and memory for patterns. *The American Journal of Psychology*, 68(2), 209–222.
- Barden, J., & Petty, R. E. (2008). The mere perception of elaboration creates attitude certainty: Exploring the thoughtfulness heuristic. *Journal of Personality and Social Psychology*, 95, 489–509.
- Barden, J., Rucker, D. D., & Petty, R. E. (2005). Saying one thing and doing another: Examining the impact of event order on hypocrisy judgments of others. *Personality and Social Psychology Bulletin*, 31(11), 1463–1474.
- Barden, J., & Tormala, Z. L. (2014). Elaboration and attitude strength: The new metacognitive perspective. *Social and Personality Psychology Compass*, 8(1), 17–29.
- Bassili, J. N. (1996). Meta-judgmental versus operative indexes of psychological attributes: The case of measures of attitude strength. *Journal of Personality and Social Psychology*, 71(4), 637–653.
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology*, 5(4), 323.
- Beale, J. M., & Keil, F. C. (1995). Categorical effects in the perception of faces. *Cognition*, 57(3), 217–239.
- Blankenship, K. L., Wegener, D. T., & Murray, R. A. (2012). Circumventing resistance: Using values to indirectly change attitudes. *Journal of Personality and Social Psychology*, 103(4), 606.
- Bornstein, M. H. (2006). Hue categorization and color naming physics to sensation to perception. *Progress in colour studies: Volume II. Psychological aspects* (pp. 35). .
- Bornstein, M. H., & Korda, N. O. (1984). Discrimination and matching within and between hues measured by reaction times: Some implications for categorical perception and levels of information processing. *Psychological Research*, 46(3), 207–222.
- Briñol, P., Petty, R. E., & Tormala, Z. L. (2006). The malleable meaning of subjective ease. *Psychological Science*, 17, 200–206.
- Cacioppo, J. T., Gardner, W. L., & Berntson, G. G. (1997). Beyond bipolar conceptualizations and measures: The case of attitudes and evaluative space. *Personality and Social Psychology Review*, 1(1), 3–25.
- Chaiken, S., & Yates, S. (1985). Affective-cognitive consistency and thought-induced attitude polarization. *Journal of Personality and Social Psychology*, 49(6), 1470.
- Clark, J. K., & Evans, A. T. (2014). Source credibility and persuasion: The role of message position in self-validation. *Personality and Social Psychology Bulletin*, 40(8), 1024–1036.
- 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.
- Davison, W. P. (1983). The third-person effect in communication. *Public Opinion Quarterly*, 47(1), 1–15.
- Douglas, K. M., & Sutton, R. M. (2004). Right about others, wrong about ourselves? Actual and perceived self-other differences in resistance to persuasion. *British Journal of Social Psychology*, 43(4), 585–603.
- Duck, J. M., Terry, D. J., & Hogg, M. A. (1995). The perceived influence of AIDS advertising: Third-person effects in the context of positive media content. *Basic and Applied Social Psychology*, 17(3), 305–325.
- Eagly, A. H., & Chaiken, S. (1993). *The psychology of attitudes*. Fort Worth, TX: Harcourt Brace.
- Ectoff, N. L., & Magee, J. J. (1992). Categorical perception of facial expressions. *Cognition*, 44(3), 227–240.
- Fabrigar, L. R., & Petty, R. E. (1999). The role of the affective and cognitive bases of attitudes in susceptibility to affectively and cognitively based persuasion. *Personality and Social Psychology Bulletin*, 25(3), 363–381.
- Faul, F., Erdfelder, E., Lang, A. G., & Buchner, A. (2007). G* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175–191.
- Fazio, R. H., & Olson, M. A. (2014). The MoDe model. *Dual-process theories of the social mind* (pp. 155). .
- Fleiss, J. L. (1971). Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76(5), 378.
- Fritz, M. S., & MacKinnon, D. P. (2007). Required sample size to detect the mediated effect. *Psychological Science*, 18(3), 233–239.
- Glasman, L. R., & Albarracín, D. (2006). Forming attitudes that predict future behavior: A meta-analysis of the attitude-behavior relation. *Psychological Bulletin*, 132(5), 778.
- Goldstone, R. L., & Hendrickson, A. T. (2010). Categorical perception. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(1), 69–78.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology*, 74(6), 1464.
- Haddock, G., Zanna, M. P., & Esses, V. M. (1993). Assessing the structure of prejudicial attitudes: The case of attitudes toward homosexuals. *Journal of Personality and Social Psychology*, 65(6), 1105.
- Harnad, S. (1987). Psychophysical and cognitive aspects of categorical perception: A critical overview. *Categorical perception: The groundwork of cognition* (pp. 1–52). Cambridge University Press.
- Hayes, A. F., & Preacher, K. J. (2014). Statistical mediation analysis with a multicategorical independent variable. *British Journal of Mathematical and Statistical Psychology*, 67(3), 451–470.
- Hoffman, M., Yoeli, E., & Nowak, M. A. (2015). Cooperate without looking: Why we care what people think and not just what they do. *Proceedings of the National Academy of Sciences*, 112(6), 1727–1732.
- Isaac, M. S., & Schindler, R. M. (2013). The top-ten effect: Consumers' subjective categorization of ranked lists. *Journal of Consumer Research*, 40(6), 1181–1202.
- Jacoby, L. L., Kelley, C., Brown, J., & Jasechko, J. (1989). Becoming famous overnight: Limits on the ability to avoid unconscious influences of the past. *Journal of Personality and Social Psychology*, 56(3), 326.
- Janis, I. L., & King, B. T. (1954). The influence of role playing on opinion change. *The Journal of Abnormal and Social Psychology*, 49, 211–218.
- Karmarkar, U. R., & Tormala, Z. L. (2010). Believe me, I have no idea what I'm talking about: The effects of source certainty on consumer involvement and persuasion. *Journal of Consumer Research*, 36(6), 1033–1049.
- Karpinski, A., & Hilton, J. L. (2001). Attitudes and the implicit association test. *Journal of Personality and Social Psychology*, 81(5), 774.
- Klein, N., & O'Brien, E. (2017). The power and limits of personal change: When a bad past does (and does not) inspire in the present. *Journal of Personality and Social Psychology*, 113(2), 210.
- Kreps, T. A., Laurin, K., & Merritt, A. C. (2017). Hypocritical flip-flop, or courageous evolution? When leaders change their moral minds. *Journal of Personality and Social Psychology*, 113(5), 730.
- Krosnick, J. A., Boninger, D. S., Chuang, Y. C., Berent, M. K., & Carnot, C. G. (1993). Attitude strength: One construct or many related constructs? *Journal of Personality and Social Psychology*, 65(6), 1132.
- Krosnick, J. A., & Petty, R. E. (1995). Attitude strength: An overview. In R. E. Petty, & J. A. Krosnick (Eds.). *Attitude strength: Antecedents and consequences* (pp. 1–24). Mahwah, NJ: Erlbaum.
- Krueger, J., & Clement, R. W. (1994). Memory-based judgments about multiple categories: A revision and extension of Tajfel's accentuation theory. *Journal of Personality and Social Psychology*, 67(1), 35.
- Labroo, A. A., & Lee, A. Y. (2006). Between two brands: A goal fluency account of brand evaluation. *Journal of Marketing Research*, 43(3), 374–385.
- Lacetera, N., Pope, D. G., & Sydnor, J. R. (2012). Heuristic thinking and limited attention in the car market. *American Economic Review*, 102(5), 2206–2236.
- Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 159–174.
- Lee, A. Y., & Aaker, J. L. (2004). Bringing the frame into focus: The influence of regulatory fit on processing fluency and persuasion. *Journal of Personality and Social Psychology*, 86, 205–218.
- Lee, A. Y., & Labroo, A. A. (2004). The effect of conceptual and perceptual fluency on brand evaluation. *Journal of Marketing Research*, 41, 151–165.
- Liberman, A. M., Harris, K. S., Hoffman, H. S., & Griffith, B. C. (1957). The discrimination of speech sounds within and across phoneme boundaries. *Journal of Experimental Psychology*, 54(5), 358.
- Luttrell, A., Petty, R. E., Briñol, P., & Wagner, B. C. (2016). Making it moral: Merely labeling an attitude as moral increases its strength. *Journal of Experimental Social Psychology*, 65, 82–93.
- Mayer, N. D., & Tormala, Z. L. (2010). “Think” versus “feel” framing effects in persuasion. *Personality and Social Psychology Bulletin*, 36(4), 443–454.
- Mishra, A., & Mishra, H. (2010). Border bias: The belief that state borders can protect against disasters. *Psychological Science*, 21(11), 1582–1586.
- Pastore, R. E. (1987). Categorical perception: Some psychophysical models. In S. Hamad (Ed.). *Categorical perception* (pp. 29–52). Cambridge, England: Cambridge University Press.
- 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.
- 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, 722–741.
- Petty, R. E., & Cacioppo, J. T. (1979). Issue involvement can increase or decrease persuasion by enhancing message-relevant cognitive responses. *Journal of Personality and Social Psychology*, 37, 1915–1926.
- Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. *Advances in Experimental Social Psychology*, 19, 123–205.
- Petty, R. E., & Krosnick, J. A. (Eds.). (1995). *Attitude strength: Antecedents and consequences*. Mahwah, NJ: Lawrence Erlbaum Associates, Inc.
- Petty, R. E., & Wegener, D. T. (1998). Matching versus mismatching attitude functions: Implications for scrutiny of persuasive messages. *Personality and Social Psychology Bulletin*, 24, 227–240.
- Pieters, R. (2017). Meaningful mediation analysis: Plausible causal inference and informative communication. *Journal of Consumer Research*, 44(3), 692–716.
- Reber, R., & Schwarz, N. (1999). Effects of perceptual fluency on judgments of truth. *Consciousness and Cognition*, 8(3), 338–342.
- Reber, R., Winkielman, P., & Schwarz, N. (1998). Effects of perceptual fluency on affective judgments. *Psychological Science*, 9(1), 45–48.
- Reich, T., & Tormala, Z. L. (2013). When contradictions foster persuasion: An attributional perspective. *Journal of Experimental Social Psychology*, 49(3), 426–439.
- Roberson, D., & Davidoff, J. (2000). The categorical perception of colors and facial expressions: The effect of verbal interference. *Memory & Cognition*, 28(6), 977–986.
- Ross, M. (1989). Relation of implicit theories to the construction of personal histories. *Psychological Review*, 96(2), 341.
- Rozin, P., & Royzman, E. B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review*, 5(4), 296–320.
- Rucker, D. D., & Petty, R. E. (2004). When resistance is futile: Consequences of failed counterarguing for attitude certainty. *Journal of Personality and Social Psychology*, 86(2), 219.

- Schwarz, N. (2004). Metacognitive experiences in consumer judgment and decision making. *Journal of Consumer Psychology, 14*, 332–348.
- Stevens, S. S. (1946). On the theory of scales of measurement. *Science, 103*, 677–680.
- Tajfel, H., & Wilkes, A. L. (1963). Classification and quantitative judgement. *British Journal of Psychology, 54*(2), 101–114.
- Tesser, A. (1978). Self-generated attitude change. *Advances in Experimental Social Psychology, 11*, 289–338.
- Thomas, M., & Morwitz, V. G. (2009). The ease-of-computation effect: The interplay of metacognitive experiences and naive theories in judgments of price differences. *Journal of Marketing Research, 46*(1), 81–91.
- Tomz, M., & Van Houweling, R. P. (2009). Candidate inconsistency and voter choice. *APSA 2009 Toronto meeting paper*.
- Tormala, Z. L., & Briñol, P. (2015). Attitude change and persuasion: Past, present, and future directions. In M. I. Norton, D. D. Rucker, & C. Lambertson (Eds.). *Cambridge handbook of consumer psychology* (pp. 29–64). Cambridge, MA: Cambridge Press.
- Tormala, Z. L., & Petty, R. E. (2002). What doesn't kill me makes me stronger: The effects of resisting persuasion on attitude certainty. *Journal of Personality and Social Psychology, 83*(6), 1298–1313.
- 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.
- Tormala, Z. L., & Rucker, D. D. (2018). Attitude certainty: Antecedents, consequences, and new directions. *Consumer Psychology Review, 1*, 72–89.
- Whittlesea, B. W., Jacoby, L. L., & Girard, K. (1990). Illusions of immediate memory: Evidence of an attributional basis for feelings of familiarity and perceptual quality. *Journal of Memory and Language, 29*(6), 716–732.
- Wicker, A. W. (1969). Attitudes versus actions: The relationship of verbal and overt behavioral responses to attitude objects. *Journal of Social Issues, 25*(4), 41–78.