Title
Facilitating Decision Making through Attribute Matching

Author
Perfecto, Hannah Jean

Publication Date
2017

Facilitating Decision Making through Attribute Matching

By

Hannah Jean Perfecto

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Business Administration

in the

Graduate Division

of the

University of California, Berkeley

Committee in Charge:

Professor Leif D. Nelson, Chair
Professor Clayton R. Critcher
Professor Ellen R. K. Evers
Professor Gabriel Lenz

Spring 2017

Abstract

Facilitating Decision Making through Attribute Matching

by

Hannah Jean Perfecto

Doctor of Philosophy in Business Administration

University of California, Berkeley

Professor Leif D. Nelson, Chair

Across 11 studies, the authors demonstrate a novel framing effect, *attribute matching*, whereby matching a salient attribute of a decision frame with that of a decision's options facilitates decision making. This attribute matching is shown to increase decision confidence and, ultimately, consensus estimates by increasing feelings of metacognitive ease. In Study 1, participants choosing the more attractive of two faces or rejecting the less attractive face reported greater confidence in and perceived consensus around their decision. Study 2 extended this finding from valence to calorie judgments, whereas Studies 3-5 extended the effect to different post-decision attitudes. Study 6 found decision ease mediates these changes in confidence and consensus estimates. Consistent with a misattribution account, when participants were warned to this external source of ease in Study 7, the effect disappeared. Studies 8-10 rule out alternative accounts, such as response substitution and language effects. The final study demonstrates attribute matching in a more realistic context. The paper concludes with a discussion of related psychological constructs as well as potential downstream consequences.

Acknowledgments

I would first and foremost like to thank my advisor, Leif D. Nelson, for his unwavering scholarly, statistical, social, and emotional support since I began my graduate studies five years ago. I would not have reached this point without him, and am truly grateful for all that he has invested in me.

I would also like to thank the rest of my committee—Clayton R. Critcher, Ellen R. K. Evers, and Gabriel Lenz—for their patience and helpful suggestions in assembling this work.

Portions of this work were presented at the Annual Meeting of the Society of Judgement and Decision Making in 2015 and at the conference for the Society of Consumer Psychology in 2016. More importantly, part of this work has been accepted for publication, co-authored by Jeff Galak, Joseph P. Simmons, and Leif D. Nelson. I would like to thank the three of them for their patience, support, and guidance through the review process. (For this reason, I maintain the use of "we" throughout the main text, regardless of a study's presence in the publication.)

Finally, I am very grateful to my fiancé, Kyle, for his constant support and understanding throughout this crazy, five-year adventure (and for his incredible cat, Skittles).

Decades of research have shown that people's preferences are often malleable. For example, we now know that people make different choices depending on whether options are framed as gains or losses (Tversky & Kahneman, 1981), or on which other options just happen to be in front of them (Simonson & Tversky, 1992; Tversky & Shafir, 1992), or on the time frame over which an attribute is described (e.g., price per year vs. price per month; Burson, Larrick, & Lynch, 2009), or on the name given to an option (Read et al., 2005), etc. Based on such findings, there is now a widespread appreciation of the power of "choice architecture," of the fact that how choice options are arrayed and described can exert a powerful influence on the decisions that people make (Kahneman & Tversky, 1984; Thaler & Sunstein, 2008).

In this paper, we describe a simple framing manipulation that affects not what people choose, but rather how they feel about their choice. Notably, we show that this framing manipulation can, by changing feelings of confidence, also influence people's beliefs about the choices of others. This is an important contribution, as practitioners are not only in the business of altering preferences; they are also in the business of altering how people feel about the preferences they already have, such as when politicians seek to strengthen the attitudes of those who are already inclined to prefer the campaign's candidate, or when marketers seek to strengthen the attitudes of those who are already inclined to prefer the company's product.

Such campaigns exist because although two prospective voters might agree in their preference, the citizen who is more confident in her preference will be more likely to vote. Indeed, people are more likely to act on behalf of preferences that are confidently held (e.g., Petty & Krosnick, 1995; Tormala & Petty, 2002), and that they think others would endorse (Goldstein, Cialdini, & Griskevicius, 2008). On the other hand, people are less likely to procrastinate in making a choice if they are confident as to which choice to make, and believe others would decide similarly (Dhar, 1997). For example, a patient deciding which of two medical procedures to undergo is more likely to make that decision quickly if she is confident about which choice to make, and if she thinks that others would make the same choice. It is therefore important to not only identify interventions that affect what people choose, but to also identify interventions that affect how people feel about those choices.

In thinking about the variables that influence decision confidence and consensus estimation, we start with the most basic of features of a decision: (1) the valence of the options, and (2) whether the decision is framed as a choice or a rejection. Option valence has already been shown to have some unexpected effects on decision satisfaction. People usually prefer to make their own decisions rather than having others decide for them (e.g., Botti, Orfali, & Iyengar, 2009; Gilovich & Medvec, 1995; Perlmuter & Monty, 1977; Stotland & Blumenthal, 1964; Taylor, Lichtman, & Wood, 1984). However, this preference is seemingly eliminated when people are forced to choose among undesirable alternatives. Choosing among negative options lowers confidence and satisfaction with the outcome (Beattie, Baron, Hershey, & Spranca, 1994; Burger, 1989), often so much so that people actively avoid making a decision at all (Botti & Iyengar, 2006). For example, research suggests that, when all of the alternatives are undesirable (e.g., two bad meal options), people are more satisfied with the outcome when someone else chooses for them than when they make the choice themselves (Botti & Iyengar, 2004). This disutility of choosing for oneself comes from the unpleasant process of focusing on the disadvantages of each outcome (see also Botti & McGill, 2006). Negative options, it seems, upend many of the benefits of choice.

Notably, all of the previous work in this area focused on decisions framed as choices – participants were asked to *choose* the option they most preferred. This brings us to the second major variable we consider: the framing of the decision. The person deciding between a chicken or steak entrée can see their decision as a choice (“I choose the chicken”) or as a rejection (“I reject the steak”). The exact same options, with the exact same outcome, might be experienced differently when framed as a choice rather than as a rejection.

Although these frames are necessarily identical in terms of outcome for binary choices, that does not mean that people think of them identically (Hubert, Neale, & Northcraft, 1987; Park, Jun, & MacInnis, 2000). Positive frames (i.e., “choose”) highlight positive attributes, whereas negative frames (i.e., “reject”) highlight negative attributes (Houston, Sherman, & Baker, 1991; Shafir, 1993). Choosing may bring about more intuitive thinking, whereas rejecting may bring about more deliberative thinking (Nagpal & Krishnamurthy, 2008; Sokolova & Krishna, 2016). More importantly, recall that previous research had suggested that a focus on negative attributes diminishes the utility of choosing between negative options (Botti & Iyengar, 2004). That is only part of the story. We propose that the match between the decision frame and the choice context (e.g., positive vs. negative options), what we will call *attribute matching*, determines the psychological response to the decision itself. That is, people will feel more confident in their decision if the frame can be changed by matching a salient attribute (valence) of the decision frame with that of the options: choose the desired or reject the undesired option.

This attribute matching, we propose, fosters a feeling of metacognitive ease (i.e., that the decision feels easy to make), which mitigates the disutility of rejecting among negative options and enhances the utility of choosing among positive options. Metacognitive ease, or fluency, has a variety of positive effects on decision making: for instance, fluent stimuli are reported as being more likable (Reber, Winkielman, & Schwarz, 1998) and more accurate (McGlone & Tofighbakhsh, 2000) than disfluent statements (see Alter & Oppenheimer, 2009, for a review). Moreover, people feel more confident in decisions that feel easy to make (Fazio & Zanna, 1978; Tormala, Petty, & Briñol, 2002). In the present paper, we investigate the role that this ease of decision making plays in determining psychological outcomes for individuals when making decisions among options where attribute matching exists versus does not exist.

We focus on a decision-making context in part because it is so broadly applicable. If a decision is experienced more positively, it should spillover into at least two critical domains: decision confidence and consensus estimates. Regarding choice confidence, decision makers are necessarily trying to identify the correct answer, but any decision will come with some sense of uncertainty (Gross, Holtz, & Miller, 1995; Kahneman, Slovic, & Tversky, 1982). In line with its ubiquity, this concept of attitude certainty has spawned a large literature examining myriad antecedents and consequences (Rucker, Tormala, Petty, & Briñol, 2014). Attitude certainty (see Tormala & Rucker (2007) for review) can even be further divided into attitude correctness and attitude clarity (Petrocelli, Tormala, & Rucker, 2007), although we focus on the broader notion of confidence in the present paper. The experience of confidence in a decision can guide how sensitive people are to other constraining information and thus their subsequent behavior (e.g., Simmons & Nelson, 2006). Accordingly, it would be both important and surprising if merely increasing the apparent match between options and frames could operate on confidence.

Regarding consensus estimation, this construct has been at the core of advances in the understanding of social judgment. Starting with initial work on the false consensus effect (Ross, Greene, & House, 1977) there has been recognition that people first look inward when asked to make judgments about others. This projection tendency, whether rational or irrational (Dawes, 1990; Krueger & Clement, 1994), springs from egocentrism. The tendency is consequential. For example, people misjudge the thirst and hunger of others depending on their own state (Van

Boven & Loewenstein, 2003) and misjudge the humor of a new joke based on their own prior exposure (Campbell et al., 2014). Those misjudgments spring from transient personal states, unaccounted for when characterizing others in a different state. Again, notably, and most interestingly, we think that this might occur as a result of merely manipulating the decision frame. People might experience the transient confidence from the match of valence and frame, and infer that their decision will be more popular than it actually is. Although previous work has found evidence for consensus estimates driving those of confidence (Horcajo, Briñol, & Petty, 2010; Petrocelli, Tormala, & Rucker, 2007; Visser & Mirabile, 2004), we present the aforementioned, more intuitive order. That is, we presume that people go through the world feeling, first and foremost, more or less confident about their decisions, and secondly (and possibly only if they are prompted) considering what percentage of others may agree with them. Regardless, in our studies, we frequently fail to attenuate our consensus effects after including confidence in the model, and reversing the order yields similarly incomplete results. The aim of this paper is not to demonstrate the order of the process, but rather to demonstrate how a novel and simple framing manipulation can affect both of these outcomes.

To our knowledge, only two published studies have come close to testing our hypothesis, but did so incompletely and with inconsistent results. Meloy and Russo (2004, Studies 2a and 2b) asked participants to either choose or reject between positive options (e.g., good employees) or negative options (e.g., bad employees), and then measured decision confidence. However, the authors did not cleanly manipulate valence, as in their studies the "positive" employee had some negative features (e.g., an employee described as "a plodder") and the "negative" option some positive features (e.g., an employee said to have periods of "above average productivity"). This muddling of valence may explain the researchers' muddled results (in fairness, these results were not of primary concern to the authors): Although people were more confident when choosing between two positive options than when rejecting between them (as we would predict), they were equally confident when choosing between two negative options as when rejecting between them (as we would not predict). Moreover, even this attenuated interaction from their Study 2a failed to emerge in their Study 2b. In our studies, we provide a much cleaner test of our hypothesis, and we consistently demonstrate the robustness of our findings across multiple stimuli and multiple studies with large samples. In addition, we propose and present evidence for a mechanism for this pattern of findings, generalize beyond simple valence matching and subjective decisions, and investigate not only reported confidence, but also perceived consensus in decision making.

Our studies investigate how attribute matching can influence confidence and consensus estimates. In Study 1, we establish evidence for the attribute matching hypothesis, showing that people are more confident in, and believe that others are more likely to agree with, choices between positive options and rejections between negative options than choices between negative options and rejections between positive options. Study 2 extends this finding to objective judgments, evaluating the calorie contents of different foods. Studies 3-5 bring more nuance to what can be matched and how we assess these effects. In Study 6, we examine whether the effects arise because the speed and ease of matched decisions inspire greater confidence. In Study 7, we show that the effect hinges on a lack of awareness, as it is eliminated when people are warned that attribute matching might influence their confidence. Studies 8-10 address two possible alternative explanations. The final study applies attribute matching to a consumer decision making context.

We report how we determined our sample size, all data exclusions (if any), all manipulations, and all measures in the respective study descriptions; we analyzed our data only after collection had finished (Simmons, Nelson, & Simonsohn, 2011). Demographic information was not always collected, but is reported whenever it was available. Sample sizes were at least 100 participants per between-subjects condition, and larger if time and resources allowed. Every study in this paper was pre-registered with the Open Science Framework. Pre-registered hypotheses, sample sizes, materials, procedures, exclusion criteria, and analysis plans, as well as full data sets for each study are linked within each study's description. For completeness, we include four studies in the Appendix originally intended for this paper, but that yielded inconclusive results.

Study 1

Each of the studies manipulating decision frame use a similar paradigm. In this study, participants were asked to make a decision between two similarly likable options. For some people this decision was expressed as a choice, whereas for others it was framed as a rejection. Additionally, we varied the valence of the pairs of options. Some pairs consisted of two desirable options and some pairs consisted of two undesirable options. In this way, we orthogonally manipulated the frame of the decision and the valence of the targets, creating matched and mismatched pairings. Moreover, our mixed design allowed all participants to experience both types of pairings. Choosing from attractive options and rejecting from unattractive options were the matched-valence trials. All participants made their decision and then reported their decision confidence and their estimate of the percentage of others who would have made the same decision. (Pre-registration and data: https://osf.io/tc7h4/ .)

Method

Participants. We recruited 2,519 participants from a private research company (M_{age} = 48.5, 51.5% female) to complete a survey about preferences. We determined this sample size in advance to be all participants in one survey session. This session consisted of multiple unrelated studies strung together, and ran until at least 1,500 participants had passed the attention check in the first study of the set. We decided in advance to analyze only the data of those who passed the check. We did not analyze the data for Study 1 until that threshold had been met and data collection had ceased.[1]

Materials, procedure, and design. Participants were asked to imagine they were selecting models for an upcoming advertising campaign. They viewed pairs of women's headshots from Pochon, Riis, Sanfey, Nystrom, and Cohen (2008), which were pretested to be attractive or unattractive. The pairs were designed such that attractiveness did not significantly differ within each pair (see Pochon et al., 2008, for more information about the pretest). Specifically, participants saw 16 pairs of women's headshots: eight pairs of attractive women and eight pairs of unattractive women, presented in a randomized order. For approximately half of participants the decision was framed as a choice (i.e., "Which woman would you choose?")

[1] The sample size of 1,500 was set for the purposes of a larger project that specifies that target for all contributing experiments. Notably, as part of that project, Study 1 was replicated in three other labs sampling from three other populations. Although the project's protocol precludes us from sharing those results at this time, we can say that the effects were highly significant and similar in magnitude to the reported study.

and for the remaining participants the decision was framed as a rejection (i.e., “Which woman would you reject?”). See Figure 1 for a sample attractive pair in the choose condition.

After each selection, participants were reminded of their answer and reported how confident they were in their decision (1 = not at all confident, 9 = extremely confident) and, as a measure of perceived consensus, what percentage of other people would make the same decision (sliding scale: 0% - 100%, with the marker starting at 50%). At the conclusion of the 16 trials, participants were given a brief attention check (Oppenheimer, Meyvis, & Davidenko, 2009) and provided their age and gender.

Figure 1. A choice screen from the Choose/Attractive condition of Study 1.

Figure 1. A choice screen from the Choose/Attractive condition of Study 1.

Results

In this study, 1,018 participants (40.4%) failed the attention check and were excluded from analyses. Although this number may seem high, it is not atypical for this more naïve participant pool, and the exclusion rule adheres to our preregistration plan.[2] Figure 2 plots the results.

Confidence. With the remaining 1,501 participants, we ran an ordinal least squares (OLS) regression. Using trial as the unit of analysis, we regressed participants’ confidence ratings on *frame* (0.5 = choose, -0.5 = reject) and *option* (0.5 = attractive faces, -0.5 = unattractive faces), and the *frame x option* interaction. We clustered standard errors at the participant level. A main effect of *option* emerged (b = 0.43, SE = 0.03 $p < .001$), as well as a

[2] Although we pre-registered to only analyze the data from those who passed the attention check, in each study the effect was very similar in size and highly significant when analyzing all participants. All data are available on the OSF pages.

main effect of *frame* ($b = 0.16$, SE = 0.06, $p = .008$). Critically, these effects were qualified by our predicted interaction ($b = 1.10$, SE = 0.06, $p < .001$): For attractive targets, participants who were choosing the better model were more confident in their decision (M = 7.04, SD = 1.04) than were the participants rejecting the worse model (M = 6.33, SD = 1.41), $t(1499) = 11.09$, $p < .001$. For unattractive targets, the effect reversed ($M_{choose} = 6.05$, $SD_{choose} = 1.34$ vs. $M_{reject} = 6.45$, $SD_{reject} = 1.24$), $t(1499) = 5.95$, $p < .001$.

Consensus. We ran the same analysis for consensus estimates. As before, although *frame* ($b = 1.96$, SE = 0.55 $p < .001$) and *option* ($b = 0.73$, SE = 0.29, $p = .013$) were both significant predictors of consensus estimates, these effects were qualified by the predicted crossover interaction ($b = 9.19$, SE = 0.58 $p < .001$): For attractive trials, participants choosing the better model gave higher consensus estimates (M = 68.4%, SD = 12.3%) than participants rejecting the worse model (M = 61.9%, SD = 12.6), $t(1499) = 10.19$, $p < .001$). However, this effect reversed for unattractive trials ($M_{choose} = 63.1\%$, $SD_{choose} = 12.0\%$ vs. $M_{reject} = 65.7\%$, $SD_{reject} = 11.1\%$; $t(1499) = 4.42$, $p < .001$). In addition, when we entered confidence measures into the model, the interaction effect for consensus estimates was reduced (from $b = 9.19$ to $b = 2.35$), consistent with mediation, $z = 4.92$, $p < .001$.

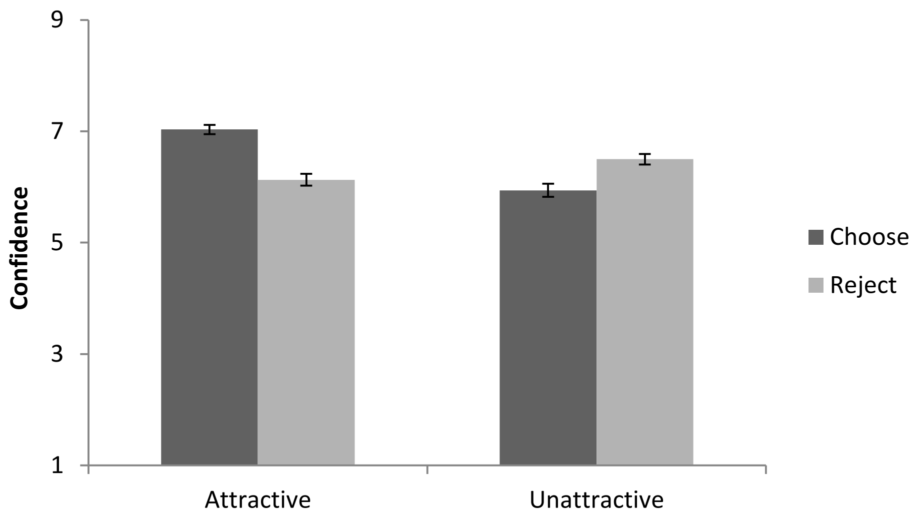

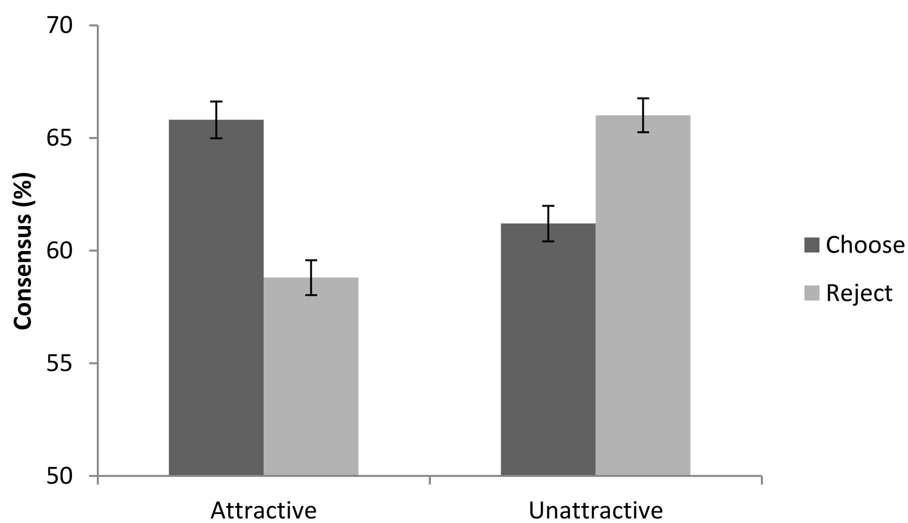

Figure 2. Matched trials (outer bars) showed higher decision confidence and perceived consensus than mismatched trials (inner bars) in Study 1.

Replication

To test whether this initial demonstration of attribute matching was replicable, we conducted an exact replication. We recruit 300 participants from Amazon Mechanical Turk and received 303 (M_{age} = 35.1, 65.6% female). The difference between our recruitment goal and final sample size (in this and subsequent Amazon Mechanical Turk studies) likely stems from a delay between the completion of the final survey and its completion code being registered online (which would close the study) or participants sharing the survey with friends for fun. In all cases, data was not analyzed until Amazon Mechanical Turk marked the study as "completed." Three participants (1.0%) failed the attention check and were excluded from analyses (Pre-registration and data: https://osf.io/bm6jr/). We ran the same analyses used in Study 1 and found an identical pattern: Confidence and consensus estimates were significantly higher when the valence of the decision frame and the options matched than when they mismatched ($b_{confidence}$ = 1.46, SE = .06; $b_{consensus}$ = 11.19, SE = 1.18, *p*s < .001). As before, when we entered confidence measures into the model, the interaction effect for consensus estimates was reduced (from *b* = 11.19 to *b* = 3.71), again consistent with mediation, *z* = 3.86, *p* < .001.

Discussion

When people were "choosing" between two attractive faces they were more confident in their decision and thought more people would agree with them than when they were "rejecting" between the exact same pair. This demonstrates a strong attribute matching effect: When people were asked to choose between options, we found that people were more confident in their decisions for attractive pairs than for unattractive pairs. However, this finding reversed under a reject frame. These matching effects on consensus estimates were then partially explained by the matching effects on confidence. Notably, the same results were obtained in a replication with a different population.

The results of Study 1 and its replication have provided preliminary evidence of attribute matching. However, these subjective judgments of faces are much more susceptible to changes in experiential information than objective judgments, which can have more declarative information to call upon (Schwarz, 1998). For a more conservative test of attribute matching, Study 2 employs a judgment task based in fact and in a domain more familiar to participants.

Study 2

Study 2 utilized a similar design as Study 1, but moved beyond valence as the attribute being matched, into a more objective domain. Instead of seeing positive and negative stimuli and being asked for their preferences, participants saw pairs of high- and low-calorie foods and were either asked which food has more calories or which food has fewer, a question with a known answer. Deciding which high-calorie food has more calories and which low-calorie food has fewer were the matched trials. (Pre-registration and data: https://osf.io/c7xez/.)

Table 1. Food pairs used in Study 2.

High-Calorie Foods	Low-Calorie Foods
Double cheeseburger vs. Medium pepperoni pizza	Baby carrot vs. Celery stick
12 buffalo chicken wings vs. Small cheese pizza	4 oz. of light orange juice vs. 4 oz. of light lemonade

Method

Participants. We received 408 participants from Amazon Mechanical Turk (M_{age} = 34.7, 62% female). We had pre-determined to recruit 400 participants in total. We did not analyze the data until that threshold had been met and data collection ceased.

Materials, procedure, and design. Participants viewed two pairs of high-calorie foods and two pairs of low-calorie foods, in a randomized order. Foods were determined to be high- or low-calorie based on a pretest, in which 51 Amazon Mechanical Turk participants were asked to estimate the caloric content of each food separately. (See Table 1 for the four food pairs.) The pairs were then constructed such that the median estimates were approximately equal. For these four food pairs, study participants were randomly assigned to determine which food has more calories or which food has fewer calories. After making each selection, participants completed the same confidence and consensus measures from the previous studies. We also asked participants how easy the decision felt to make. At the conclusion of the study, participants received a brief attention check ("Which food has more [fewer] calories? 3 grapes vs. Grilled cheese sandwich"), reported if they were vegetarian as well as if they were on a diet, and provided their age and gender.

Results and Discussion

Eleven participants (2.7%) failed the attention check and were excluded from analyses.

Ease. With the remaining participants, we ran an OLS regression. Using trial as the unit of analysis, we regressed participants' ratings of decision ease on *frame* (-0.5 = more, 0.5 = less) and *option* (0.5 = high-calorie foods, -0.5 = low-calorie foods), and the *frame x option* interaction. We clustered standard errors at the participant level. We found no effect of *frame* (b = -0.07, SE = 0.10, p = .468), although a significant effect of *option* (b = 0.42, SE = 0.07, p < .001) did emerge. More importantly, our predicted interaction was also significant (b = 0.75, SE = 0.14, p < .001): on high-calorie trials, participants who were asked which food has more calories reported the decision was easier to make (M = 4.55, SD = 1.16) than did participants who were asked which has fewer (M = 4.11, SD = 1.1), t(395) = 3.91, p < .001. However, this effect reversed when participants were asked about low-calorie foods (M_{more} = 4.60, SD_{more} = 1.37 vs. M_{fewer} = 4.91, SD_{fewer} = 1.19; t(395) = 2.35, p = .019).

Confidence. We next ran the same analysis for reported confidence. Again, we found no effect of *frame* (b = 0.07, SE = 0.14, p = .618), but did find a significant effect of *option* (b = 0.71, SE = 0.10, p < .001). More importantly, our predicted interaction was significant, as well (b = 0.90, SE = 0.19, p < .001): on high-calorie trials, participants who were asked which food has more calories gave higher confidence estimates (M = 5.39, SD = 1.68) than did participants who were asked which has fewer (M = 5.01, SD = 1.53), t(395) = 2.37, p = .018. However, this effect reversed for low-calorie trials (M_{more} = 5.65, SD_{more} = 1.86 vs. M_{fewer} = 6.17, SD_{fewer} = 1.58; t(395) = 3.00, p = .003). In addition, when we added decision ease to the model, the interaction effect for reported confidence was reduced to non-significance (from b = 0.90 to b = 0.12), consistent with mediation, z = 5.19, p < .001

Consensus. Finally, we fit the same model for consensus estimates. Again, we found no effect of *frame* (b = -0.39, SE = 1.05, p = .713), but did find a significant effect of *option* (b = 5.85, SE = .86, p < .001). More importantly, the predicted crossover interaction was again highly significant (b = 7.66, SE = 1.73, p < .001): On high-calorie trials, participants who were asked which food has more calories estimated greater consensus around their answer (M = 62.0%, SD = 13.8) than did participants who were asked which has fewer (M = 57.8%, SD = 12.9), t(395) =

3.15, p = .002. However, this effect reversed for low-calorie trials (M_{more} = 64.1%, SD_{more} = 14.5 vs. M_{fewer} = 67.5%, SD_{fewer} = 13.2; t(395) = 2.47, p = .014). In addition, when we entered reported confidence into the model, the interaction effect for consensus estimates was reduced (from b = 7.66 to b = 3.63), consistent with mediation, z = 5.07, p < .001.

Study 2 shows that the effects of attribute matching are not limited to judgments of subjective preference, but emerge even when people are judging stimuli on objective dimensions. To our knowledge, this is the first demonstration of a matching effect on objective judgments. Even when asked to make these objective judgments of caloric content, participants' confidence and perceived consensus were higher when the decision frame and the food pairs matched on their salient attribute. Moreover, participants' perceptions of how easy the decision felt to make showed this pattern as well, suggesting a similar process as in previous studies despite the new domain.

Study 3

In Study 3, although we return to preferences, we sought to test whether this effect is larger among stimuli that are more extremely (vs. moderately) valenced. As in Study 1, we varied the valence of the decision frame (choose vs. reject) and the valence of the options (positive vs. negative). Additionally, to test whether attribute matching would emerge with less extreme stimuli, we included four levels of valence: extremely positive, slightly positive, slightly negative, and extremely negative (see Table 1 for exact stimuli). We predicted that attribute matching would still emerge for the less extreme stimuli, albeit in smaller magnitudes. (Pre-registration and data: https://osf.io/expw5/.)

Method

Participants. We decided in advance to recruit 300 participants from Amazon Mechanical Turk, and received 301 (M_{age} = 36.1, 55.0% female).

Materials, procedure, design. Participants viewed 20 pairs of words in a randomized order. The words were chosen based on valence ratings provided by Bellezza, Greenwald, and Banaji (1986) and paired such that valence did not significantly differ within each pair. The 20 trials consisted of five pairs of extremely positive, slightly positive, slightly negative, and extremely negative words each (see Table 2). As in Study 1, half of participants were asked to indicate their preference by choosing the word they preferred, and half were asked to do so by rejecting the word they did not prefer. (Readers interested in the [non-differing] choice share within each pair are referred to the Appendix.) After making each choice, participants completed the same confidence and consensus measures from Study 1. At the conclusion of the study, participants received the same attention check as in Study 1 and provided their age and gender.

Table 2. Word pairs used in Studies 3, 6, and 7 (**indicates use in Studies 6-7).

Very Negative**	**Slightly Negative**	**Slightly Positive**	**Very Positive****
murderer vs. tumor	thorn vs. jealousy	circus vs. world	joy vs. kiss
poison vs. slaughter	snob vs. beggar	fur vs. privacy	pleasure vs. vacation
war vs. maggot	useless vs. wasp	knowledge vs. learn	family vs. laughter
cancer vs. funeral	rage vs. stress	water vs. employment	paradise vs. sunrise
lice vs. suicide	putrid vs. stupid	earth vs. improve	romantic vs. love

Results and Discussion

Fourteen participants (4.7%) failed the attention check and were excluded from analyses. Figure 3 plots the results.

Confidence. Using trial as the unit of analysis, we regressed confidence ratings on *frame*, *option*, *extremity* (-0.5 = moderate, 0.5 = extreme), all possible two-way interactions, and the three-way *frame x option x extremity* interaction. We clustered standard errors at the participant level. Although *frame* (b = 0.17, SE = 0.11, p = .135) and *option* (b = -0.07, SE = 0.06, p = .217) were not significant predictors of confidence, our predicted *frame x option* interaction was highly significant (b = 2.10, SE = 0.12, p < .001): On positive trials, participants choosing the better word were more confident in their choice (M = 6.92, SD = 1.04) than participants rejecting the worse word (M = 6.04, SD = 1.20), $t(285) = 6.63$, p < .001, an effect that reversed for negative trials (M_{choose} = 5.79, SD_{choose} = 1.35 vs. M_{reject} = 7.02, SD_{reject} = 0.89; $t(285) = 9.09$, p < .001). Supporting the idea that attribute extremity moderates the matching effect, the three-way interaction among *frame*, *option*, and *extremity* was highly significant, b = 1.37, SE = 0.17, p < .001. Attribute matching occurred even for the slightly positive and slightly negative word pairs, but was less pronounced, $b_{interaction}$ = 1.38, SE = 0.13, p < .001.

Consensus. We then repeated this analysis with consensus estimates. The effect of *frame* was not significant (b = -0.47, SE = .96, p = .620), whereas the effect of *option* was (b = -1.89, SE = 0.54, p < .001). Critically, however, our predicted interaction was highly significant (b = 13.72, SE = 1.09, p < .001): on positive trials, participants choosing the better word gave higher consensus estimates (M = 65.0%, SD = 10.1%) than participants rejecting the worse word (M = 57.6%, SD = 8.8%), $t(285) = 6.61$, p < .001. However, this effect reversed for negative trials (M_{choose} = 56.1%, SD_{choose} = 10.3% vs. M_{reject} = 62.5%, SD_{reject} = 9.3%; $t(285) = 5.49$, p < .001). As with confidence ratings, the three-way interaction was also significant (b = 11.86, SE = 1.71, p < .001), indicating that although the *frame x option* interaction was significant even for the moderately valenced trials ($b_{interaction}$ = 7.61, SE = 1.30, p < .001), it was smaller than it was for the extremely valenced trials ($b_{interaction}$ = 19.66, SE = 1.48, p < .001). As in Study 1, when we entered both the confidence and consensus measures into the model, the frame/option interaction effect for consensus estimates was reduced (to b = 3.65 from b = 13.74), consistent with mediation, z = 3.72, p < .001.

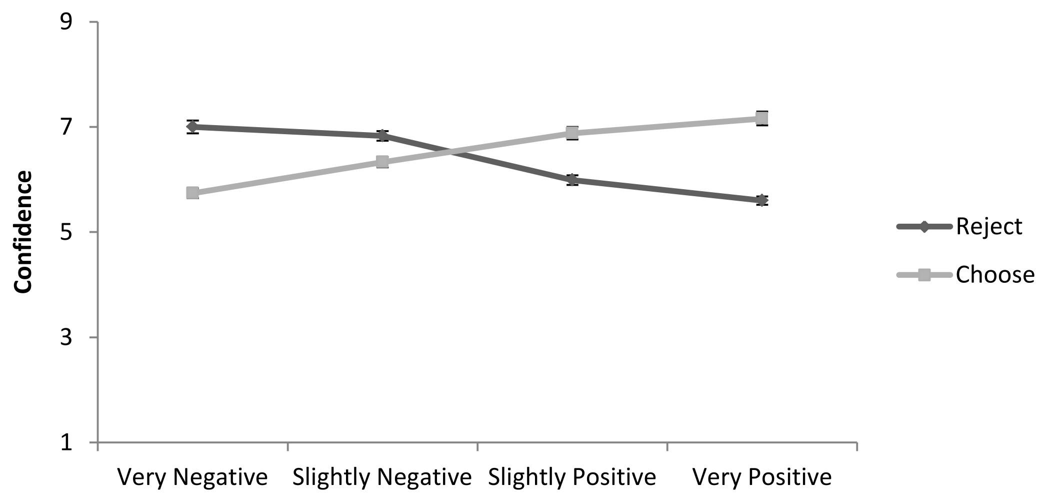

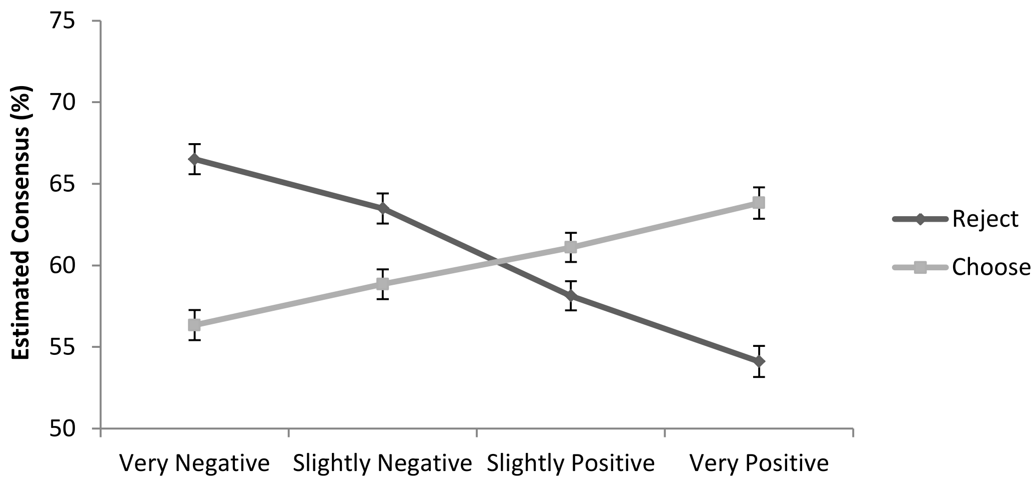

Figure 3. The size of the attribute matching effect varies with attribute intensity in Study 3.

Study 3 shows that attribute matching occurs not just when the attribute is at its extremes; even the only slightly positive and slightly negative word pairs significantly showed attribute matching, with the size of the effect varying with attribute intensity. In the next two studies, we turn to a possible mechanism for attribute matching: decision ease.

Study 4

Up until now, we have only matched salient attributes of the decision frame and the decision's options. However, in these studies, we have kept the salient attribute of our follow-up questions constant. That is, although we varied the valence of the decision frame and options (as in Studies 1 and 3), our questions about how participants felt about their decisions have always been framed positively (e.g., what percentage of other participants would have made the *same* choice? [emphasis added]). Therefore, in this exploratory study, we manipulate the framing of these follow-up questions. It could be that including a matching frame for these questions further increases decision confidence. Alternatively, it could be that participants' feelings about the decision emerge at the time of deciding, and therefore cannot be so easily influenced by subsequent framing effects. Pre-registration and data: https://osf.io/vf5h6/).

Method

Participants. We recruited 1,626 participants (M_{age} = 48.3, 48.3% female) through the same private research firm from Study 1. As in Study 1, we determined this sample size in advance to be all participants in one survey session. This session consisted of multiple unrelated studies strung together, and ran until at least 1,500 participants had passed the attention check in the first study of the set. We decided in advance to analyze only the data of those who passed the check. We did not analyze the data for Study 4 until that threshold had been met and data collection had ceased.

Materials, procedure, and design. Participants viewed the five extremely positive and five extremely negative word pairs from Study 3, in a randomized order. As in Study 3, participants were randomly assigned either to choose the word they prefer or reject the word they did not prefer. Participants were also randomly assigned to one of two follow-up frames: In the positive (negative) condition, participants were asked to report how easy (difficult) the decision felt to make, how certain (uncertain) that they chose the better word (on choose/positive trials), and what percentage of other respondents would make the same (a different) decision. After the final trial, participants received the same attention check as in previous studies and reported their age and gender.

Results and Discussion

In this study, 630 participants (37.1%) failed the attention check and were excluded from analyses. Using trial as the unit of analysis, for each outcome, we regressed the variable of interest on *frame* (-0.5 = choose, 0.5 = reject), *option* (-0.5 = negative, 0.5 = positive), *outcome* (-0.5 = negative, 0.5 = positive), all possible two-way interactions, and the three-way *frame X option X outcome* interaction. We clustered standard errors at the participant level. Each negatively-framed outcome was reverse-scored.

Ease/Difficulty. With the participants who passed the attention check, we found significant effects of *frame* (b = -0.29, SE = 0.06, $p < .001$) and *option* (b = 0.21, SE = 0.04, $p < .001$), as well as our predicted interaction between the two (b = -2.51, SE = 0.08, $p < .001$). On positive trials, participants choosing the better word reported the decision was easier to make (M = 5.98, SD = 0.88) than participants rejecting the worse word (M = 4.37, SD = 1.35), $t(995)$ = 21.28, $p < .001$, an effect that reversed for negative trials (M_{choose} = 4.46, SD_{choose} = 1.33 vs. M_{reject} = 5.42, SD_{reject} = 1.11; $t(995)$ = 12.47, $p < .001$). No other effects were significant, including the three-way interaction (b = -0.18, SE = 0.16, p = .255).

Certainty/Uncertainty. We next ran the same analyses on reported certainty. Here, *frame* remained a significant predictor (b = -0.44, SE = 0.09, $p < .001$), but *option* did not (b = 0.00, SE = 0.05, p = .998); importantly, their predicted interaction was still significant (b = -2.31, SE = 0.09, $p < .001$). On positive trials, participants choosing the better word reported greater certainty in their decision (M = 7.36, SD = 1.30) than participants rejecting the worse word (M = 5.76, SD = 1.77), $t(995)$ = 16.22, $p < .001$, an effect that reversed for negative trials (M_{choose} = 6.21, SD_{choose} = 1.59 vs. M_{reject} = 6.92, SD_{reject} = 1.49; $t(995)$ = 7.23, $p < .001$). However, *outcome* was also a significant predictor (b = -0.18, SE = 0.09, p = .036), as was the three-way interaction (b = -0.84, SE = 0.18, $p < .001$). Probing this interaction further, the *frame x option* interaction (i.e., our matching effect) was larger with the positive (b = -2.57, SE = .10, $p < .001$) than in the negative (b = -1.73, SE = 0.11, $p < .001$) frame. (See Figure 4 for full results.) No other effects were significant.

Agreement/Disagreement. We next ran the same analyses on estimated consensus (or lack thereof). Here, *frame* (b = -1.58, SE = 0.90, p = .048) and *option* (b = -2.90, SE = 0.46 $p < .001$) were both significant predictors, as was their predicted interaction (b = -9.40, SE = 0.92, $p < .001$). On positive trials, participants choosing the better word estimated greater consensus around their decision (M = 61.1%, SD = 16.4%) than participants rejecting the worse word (M = 55.1%, SD = 13.3%), $t(995)$ = 6.35, $p < .001$, an effect that reversed for negative trials (M_{choose} = 59.4%, SD_{choose} = 14.5% vs. M_{reject} = 62.9%, SD_{reject} = 17.6%; $t(995)$ = 3.34, $p < .001$). However, *outcome* was also a significant predictor (b = 10.16, SE = 0.80, $p < .001$), as was the three-way interaction (b = -18.45, SE = 1.84, $p < .001$). Probing this interaction further, as with certainty, the *frame x option* interaction was larger with the positive (b = -19.29, SE = 1.03, $p < .001$) than in the negative (b = 1.19, SE = 1.04, p = .255) frame. (Again, see Figure 4 for full results.) No other effects were significant.

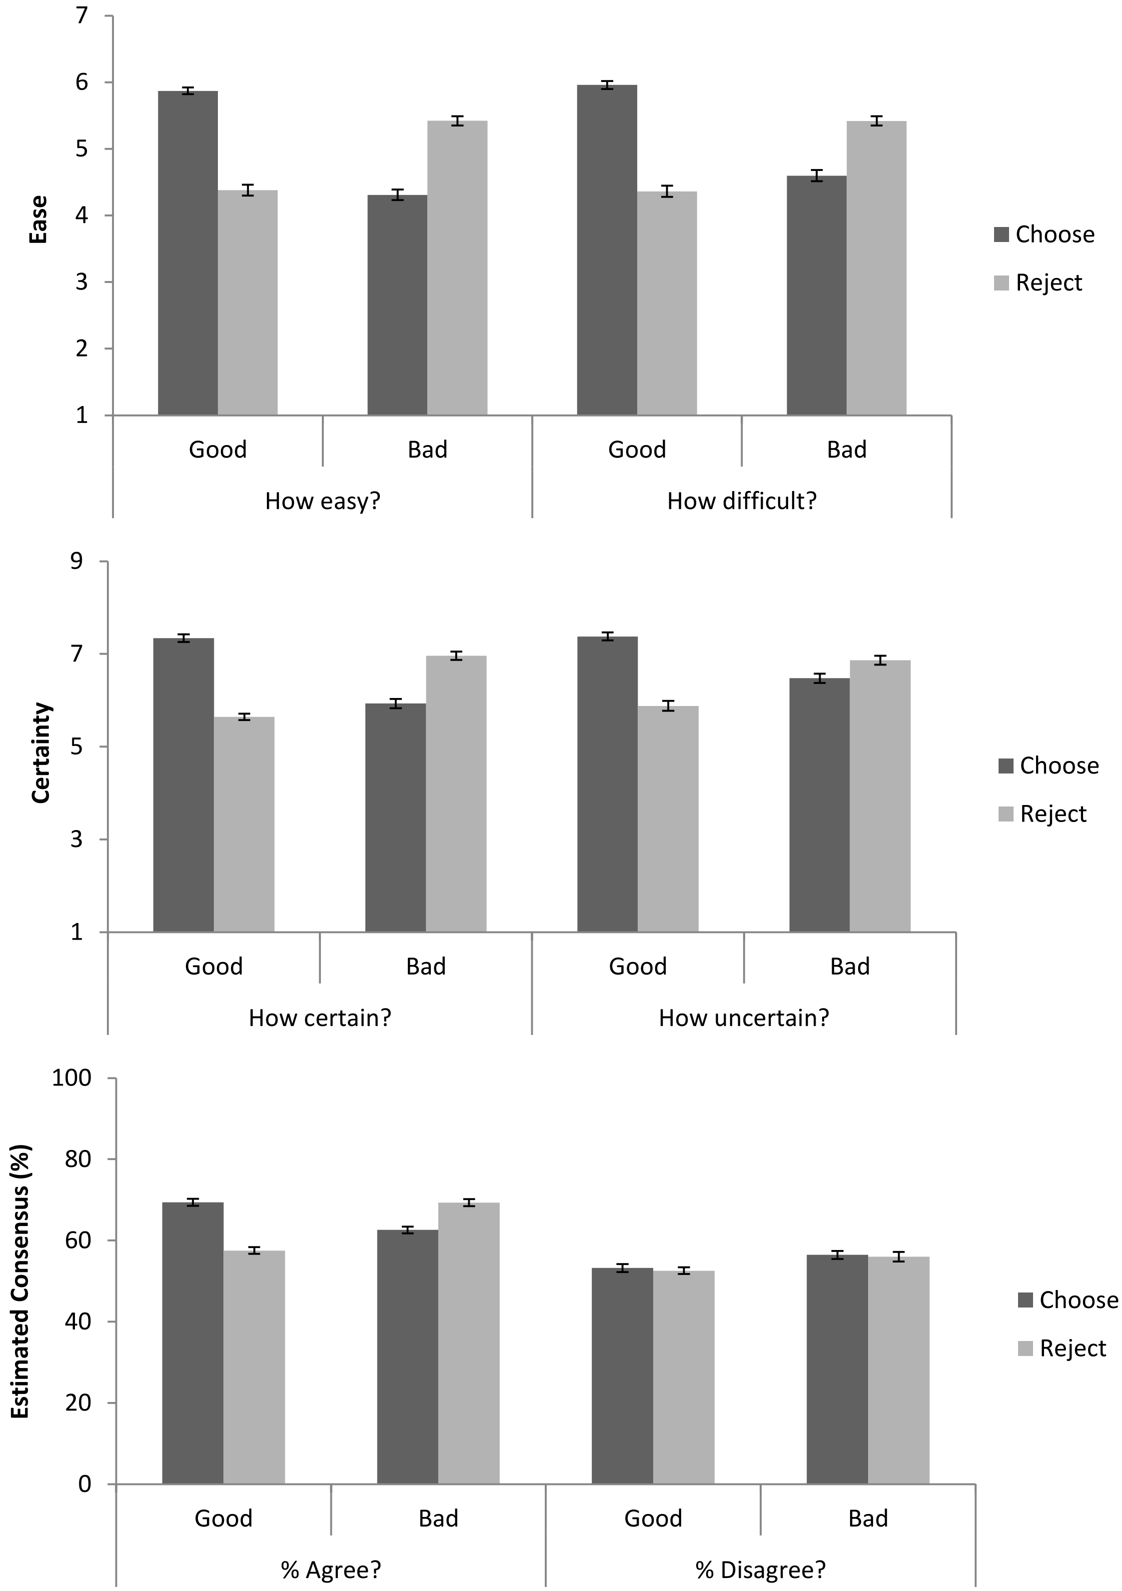

Figure 4. Varying the valence of the follow-up questions does not reliably enhance or mitigate matching effects in Study 4.

In total, then, the valence of the questions about one's decision making did not seem to enhance matching effects. Although two of the three outcomes (certainty and consensus) were affected by this manipulation, it was almost exclusively in the form of lowering all responses, rather than selectively increasing matched responses and lowering mismatched ones. Our reasoning for featuring only positively-framed questions initially was their greater frequency in everyday speech—perhaps this overall reduction stems from relative unfamiliarity with the negative frames.

Study 5

Having examined the valence of the outcomes in Study 4, the goal of Study 5 was to investigate the positive outcomes we had been using more closely. In three of the four previous studies (and in the majority of later ones), we ask participants for their confidence or certainty in their decision, as well as what percentage of other participants would make the same decision. However, one could argue that these do not capture the bulk of one's post-decision feelings. Hence, in Study 5, we ask participants to tell us about how they felt about their decision on a variety of measures. We predicted that there is nothing special about confidence and consensus ratings, and that attribute matching will emerge on all of our follow-up questions about how they felt about their decision. Pre-registration and data: https://osf.io/wupe8/).

Method

Participants. We recruited 304 participants from Amazon Mechanical Turk. We had pre-determined to recruit 300 participants in total.

Materials, procedure, and design. Participants were asked to imagine they were invited to two dinners at two friends' houses, and were asked to help their friend decide what to make. They then viewed a pair of appealing foods (T-bone steak and fried chicken) and a pair of unappealing foods (cow's tongue and tripe [parenthetically defined for participants as sheep's stomach) in a randomized order. These foods were pretested by 52 Amazon Mechanical Turk workers to be un/appealing and equated in each pair. As in previous studies, some participants were asked to choose the food they would like to have; others were asked to reject the food they would not like to have. Importantly, after each decision, participants were asked not only about their confidence in and perceived consensus around their decision, but also their certainty that they selected the better option, their happiness with the decision, their satisfaction with their decision, how enjoyable the decision was to make, how likely they would be to make the same selection (given the same options) in the future, and how unpleasant the decision was to make. Each question ranged from 1 – Not at all [X] to 9 – Very [X]. These outcomes were presented in a randomized order, with one question per page. After completing both trials, participants completed the same attention check as previous studies and reported whether they were vegetarian or on a diet.

Results and Discussion

Forty participants (13.2%) failed the attention check and were excluded from analyses. For the sake of readability, rather than report the full model results for each outcome separately, we direct the reader to the accompanying figure.

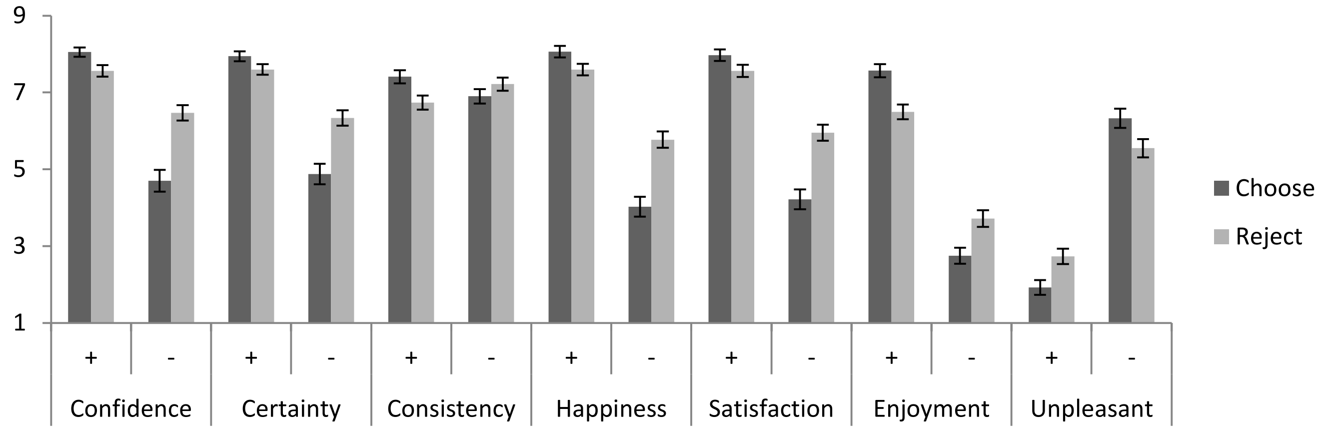

Figure 5. Attribute matching emerges across a variety of outcomes in Study 5.

Although the main effect of option valence is unusually pronounced with these stimuli, the pattern is remarkably similar both comparing outcomes to outcomes (Cronbach's α = .87) and to other studies thus far. With the exception of consistency and its marginal simple effect for unappealing foods (p = .077), participants reported feeling better about their decision when choosing from appealing foods or rejecting from unappealing foods on every dimension we included. Therefore, we do not believe that the attribute matching effect is somehow limited or unique to confidence and consensus reports.

Having learned a bit more about different domains of judgments attribute matching could emerge in in Studies 1 and 2, and about the different types of outcomes it can influence (Studies 3-5), we turn to what may be driving these changes in participants' attitudes towards their decisions in Studies 6 and 7.

Study 6

Study 6 mirrored Study 3 except for two key differences. First, Study 6 only used the extremely positive and negative words, and second, it included both direct (seven-point scale) and indirect (decision time) measures of decision ease. Decisions on matching trials should feel easier to make, and be made faster, than those for mismatched trials. Furthermore, we predicted the standard matching effect from Study 3 would replicate, but, more importantly, that it would be mediated by our direct measure of decision ease. (Because of the unreliability of response time data [Fazio, 1990; Evans, Dillon, & Rand, 2015], we pre-registered our response time variable as exploratory. Pre-registration and data: https://osf.io/jk3mu/)

Method

Participants. We recruited 421 participants from Amazon Mechanical Turk (M_{age} = 31.1, 57.2% female). We had pre-determined to recruit 400 participants in total.

Materials, procedure, and design. Participants viewed the five extremely positive and five extremely negative word pairs from Study 3, in a randomized order. As in Study 3, participants were asked either to choose the word they prefer or reject the word they did not prefer. We surreptitiously recorded the amount of time participants spent on this decision page. Next, all participants were asked how easy making their decision felt (1 – Very difficult, 9 – Very easy), before also responding to the same confidence and consensus measures we used in the previous studies. After the last trial, participants received the same attention check as in previous studies and reported their age and gender.

Results and Discussion

Sixty (14.3%) participants failed the attention check and were excluded from analyses.

Ease. With the remaining participants, we ran the same analysis as in Study 1 on the ease measure (*frame*: -0.5 = choose, 0.5 = reject; *option*: -0.5 = negative, 0.5 = positive), clustering standard errors at the participant level and using trial as the level of analysis.

Response Time. Per our pre-registration, due to the nature of response times, we first excluded any responses less than 200ms and then log-transformed the resulting data before conducting our analyses (Wheelan, 2008); however, for ease of understanding, we report the raw means in text[3]. We then fit the model for our transformed response time variable. Here, *frame* was a significant predictor of response time (b = 0.19, SE = 0.03, $p < .001$) and *option* was not (b = -0.02, SE = 0.01, p = .117). More importantly, however, our predicted interaction again obtained (b = 0.28, SE = 0.27, $p < .001$). In line with our hypothesis, on positive trials, participants who chose the words they preferred chose faster (M = 5.39s, SD = 2.82s) than did participants who rejected the words they did not prefer (M = 7.46s, SD = 3.27), $t(357) = 8.52$, $p < .001$. This difference was smaller on negative trials (M_{choose} = 6.56s, SD_{choose} = 4.04 vs. M_{reject} = 7.62s, SD_{reject} = 12.55s), $t(367) = 1.28$, p = .203.

Confidence. We next ran the same analyses on reported confidence. Here, neither *frame* (b = -0.14, SE = 0.11, p = .231) nor *option* (b = 0.11, SE = 0.08, p = .182) were significant predictors of confidence; however, our predicted interaction was (b = -2.61, SE = 0.17, $p < .001$): on positive trials, participants who chose the words they preferred reported higher confidence in making their choice (M = 7.50, SD = 0.97) than did participants who rejected the words they did not prefer (M = 6.06, SD = 1.58), $t(359) = 10.50$, $p < .001$, an effect that reversed for negative trials (M_{choose} = 6.08, SD_{choose} = 1.55 vs. M_{reject} = 7.25, SD_{reject} = 1.14; $t(359) = 8.15$, $p < .001$). When we entered ease into the model for confidence, the frame/option interaction was reduced (from b = -2.61 to b = -0.18), suggesting mediation, z = 1.99, p = .047.

Consensus. Finally, we repeated this analysis with consensus estimates: *Frame* was not a significant predictor of consensus estimates (b = -1.08, SE = 0.92, p = .243), *option* was (b = -3.65, SE = 0.67, $p < .001$), but, critically, our predicted interaction obtained (b = -14.30, SE = 1.34, $p < .001$): on positive trials, participants who chose the words they preferred gave higher consensus estimates in making their choice (M = 64.4%, SD = 11.6%) than did participants who rejected the words they did not prefer (M = 56.1%, 9.3%), $t(359) = 7.41$, $p < .001$, an effect that reversed for negative trials (M_{choose} = 60.9%, SD_{choose} = 10.7% vs. M_{reject} = 66.9%, SD_{reject} = 11.7; $t(359) = 5.16$, $p < .001$). When we entered the measure of confidence into the model predicting consensus, the interaction effect was reduced (from b = -14.30 to b = -4.12), suggesting mediation, z = 4.01, p = .005.

Study 6 presents preliminary evidence for decision ease as a mechanism behind attribute matching effects on confidence. In addition to replicating the attribute matching effect from Studies 1 and 2, participants' reports of how easy making the decisions felt, as well as how quickly they made those decisions, mediated their reported confidence. Along these lines, we hypothesize that participants misattribute this change in decision ease to their confidence in their preferences, instead of to our manipulations. However, rather than affecting only *subjective* ease, it could be that our manipulations affect how objectively easy the decisions were to make.

---

[3] Using untransformed response times yields equivalent effect sizes and significance levels. We report transformed statistical tests in keeping with convention and our pre-registration.

To disentangle these two possibilities in Study 7, we prompt some participants to consider this irrelevant source of experiential information, without stating its direction of influence. If our effect has been emerging from misattribution, this prompt should attenuate the effect; if, instead, the effect exists only because matched decisions are inherently easier, then this prompting should have no effect. This approach also allows us to implicate decision ease in our model more directly, beyond the mere correlational evidence from Study 6: Participants cued to discount their ease of decision making failing to report the usual changes in confidence and consensus estimates would suggest that decision ease is in fact involved in this process.

Study 7

The purpose of Study 7 was to more directly test the hypothesis that the increase in confidence and consensus estimates on matched trials comes from a misattribution of decision ease from attribute matching, and not from the decision becoming objectively easier to make. Study 7 used the same materials and procedure as Study 6, but without the additional measure of ease. (We chose not to measure response time in this study, given the ambiguity in the literature and the results of Study 6.) Instead, we introduced a third factor: some participants were warned that the valence of the frame and the options may have made the decision feel easier prior to their reports of confidence and consensus. We predicted that participants who did not receive this notice would show the attribute matching effect, whereas those who did would properly attribute their increase in ease to our manipulations and show no effects. (Pre-registration and data: https://osf.io/gdrvs/.)

Method

Participants. We recruited 502 participants from Amazon Mechanical Turk (M_{age} = 30.8, 49.4% female). We had pre-determined to recruit 500 participants in total. We did not analyze the data until that threshold had been met and data collection ceased.

Materials, procedure, and design. As in Study 6, participants viewed five extremely positive and five extremely negative word pairs in a randomized order, and were randomly assigned either to choose the word from each pair that they prefer or to reject the word they do not prefer. As in previous studies, all participants reported their confidence and perceived consensus. However, after making their selection but before reporting their confidence, consistent with the approach taken by Cesario, Grant, and Higgins (2004), half of participants read: "Before continuing, please consider the following: Past research suggests that phrasing a decision positively [negatively] could affect how easy your decision seems, *depending on* the positivity or negativity of the options." Participants in these warned conditions saw the message on every trial, and could not proceed to the next page for two seconds. At the conclusion of the study, participants received the same attention check as in previous studies and provided their age and gender.

Results and Discussion

Forty-four participants (8.8%) failed the attention check and were excluded from analyses. Figure 6 plots the results.

Confidence. Using trial as the unit of analysis, we regressed confidence ratings on *frame*, *option*, *warning* (-0.5 = unwarned, 0.5 = warned), all possible two-way interactions, and the three-way *frame x option x warning* interaction. We clustered standard errors at the participant level. *Frame* (b = -0.28, SE = 0.09, p = .002) and *option* (b = 0.34, SE = 0.06, p < .001) were both significant predictors of reported confidence. As in previous studies, these effects were qualified by the predicted interaction (b = -0.94, SE = 0.12, p < .001), but, more importantly, the predicted three-way interaction was also significant, *b* = 2.76, SE = 0.23, $p < .001$. To probe the nature of this interaction further, we discuss the unwarned and warned conditions separately.

In the unwarned conditions, without the warning manipulation, attribute matching replicated as predicted. The predicted *frame x option* interaction was highly significant, *b* = 2.32, SE = 0.18, $p < .001$: on positive trials, participants reported greater confidence when choosing (M = 7.43, SD = 0.99) than rejecting (M = 5.99, SD = 1.54), $t(244) = 8.61$, $p < .001$; whereas on negative trials, participants reported greater confidence when rejecting (M = 7.19, SD = 1.06) than choosing (M = 6.32, SD = 1.22), $t(244) = 5.89$, $p < .001$. In the warned conditions, however, the results were very different. Although the *frame x option* interaction was still significant (*b* = 0.44, SE = .15, $p = .003$), for positive trials participants reported similar levels of confidence whether they were choosing (M = 7.11, SD = 0.92) or rejecting (M = 7.04, SD = 0.99), $t(254) = 0.63$, $p = .530$. On negative trials, an unpredicted significant difference did emerge, but in the direction opposite of what was observed in the unwarned condition ($M_{choose} = 6.52$, $SD_{choose} = 1.21$ vs. $M_{reject} = 6.00$, $SD_{reject} = 1.29$; $t(254) = 3.53$, $p = .001$).

Consensus. We then fit the same model for consensus estimates. Neither f*rame* (*b* = -1.65, SE = 0.87, $p = .060$) nor option valence (*b* = 0.89, SE = 0.61, $p = .146$) significantly predicted estimated consensus. As in previous studies, however, the predicted *frame x option* interaction was highly significant (*b* = -3.84, SE = 1.22, $p = .002$), and more importantly, the predicted three-way interaction was also significant, *b* = 22.93, SE = 2.44, $p < .001$. To probe the nature of this interaction further, we again discuss the unwarned and warned conditions separately.

In the unwarned conditions, without the warning, attribute matching replicated as predicted. The predicted *frame x option* interaction was highly significant, *b* = -15.30, SE = 1.88, $p < .001$: on positive trials, participants estimated higher consensus when choosing (M = 65.6%, SD = 12.4%) than rejecting (M = 55.4%, SD = 11.4%), $t(244) = 6.75$, $p < .001$; whereas on negative trials, participants estimated higher consensus when rejecting (M =67.4%, SD = 11.3%) than choosing (M = 62.5%, SD = 12.9%), $t(244) = 2.97$, $p = .003$. In the warned conditions, however, the results were very different. The *frame x option* interaction was still significant (*b* = 7.63, SE = 1.55, $p < .001$), but, as with confidence, the unpredicted opposite pattern emerged: On positive trials, participants estimated a smaller consensus when choosing (M = 63.5%, SD = 9.14%) than rejecting (M = 67.2%, SD = 11.3%), $t(254) = 2.89$, $p = .004$. On negative trials, we again observed the opposite of the unwarned conditions' results ($M_{choose} = 60.6\%$, $SD_{choose} = 10$ 6% vs. $M_{reject} = 57.1\%$, $SD_{reject} = 10.2\%$; $t(254) = 2.72$, $p = .007$).

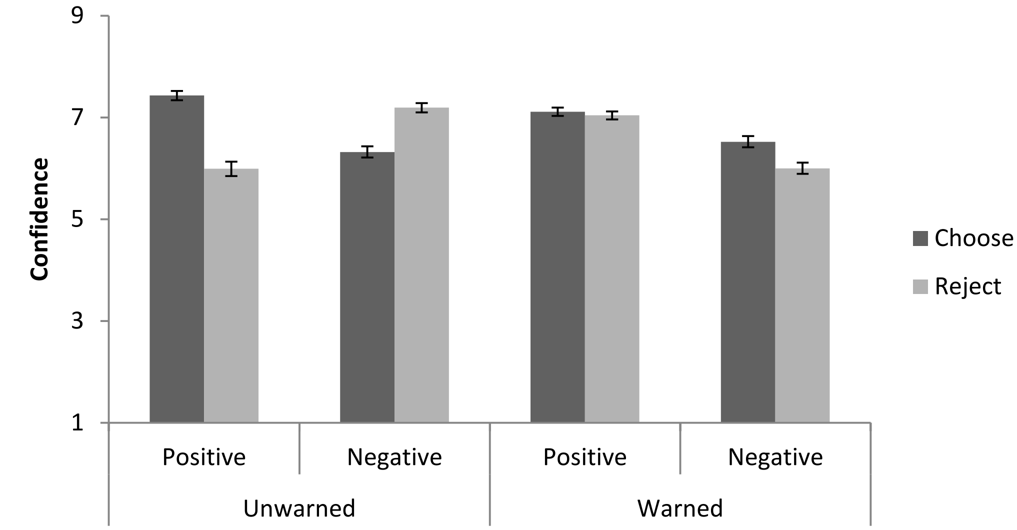

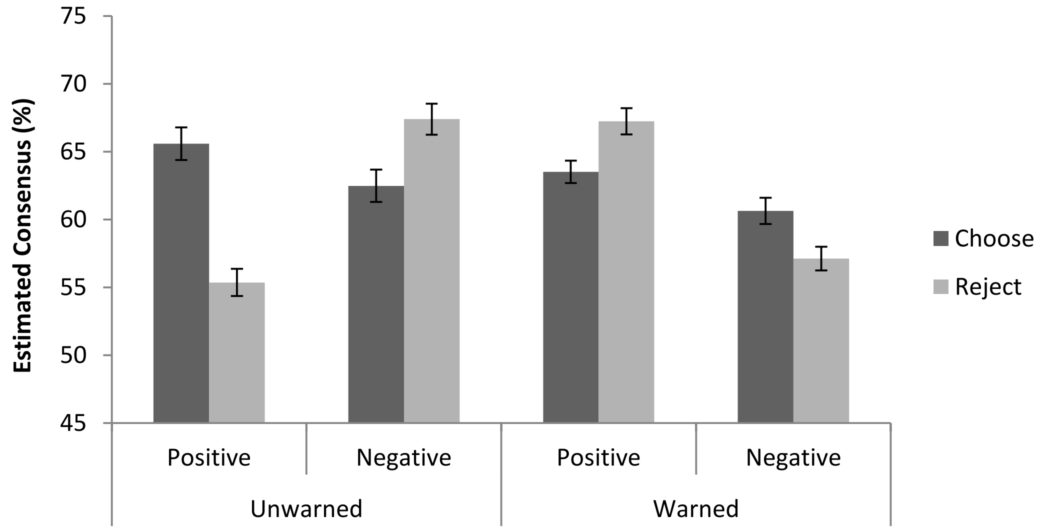

Figure 6. Attribute matching fails to emerge after the warning in Study 7.

Replication

Because of the theoretical importance of Study 7 to the present paper, we ran a direct replication with the exact same materials to verify its reliability (pre-registration and data: https://osf.io/4tzig/). We recruited 502 participants from Amazon Mechanical Turk; forty-four (8.8%) failed the attention check and were excluded from analyses. We fit the same model as before: the three-way interaction again emerged for both confidence (b = 2.76, SE = 0.23, p < .001) and consensus (b = 22.93, SE = 2.44, p < .001). The unwarned conditions showed the usual *frame x option* interaction ($b_{confidence}$ = -2.32, SE = 0.18, p < .001; $b_{consensus}$ = -15.30, SE = 1.89, p < .001). However, the warned conditions did not ($b_{confidence}$ = .44, SE = 0.15, p = .003; $b_{consensus}$ = 7.63, SE = 1.55, p < .001), showing instead the pattern from the previous study.

Discussion

Study 7 offers further, more direct, support for decision ease as a mechanism behind attribute matching's effects on confidence and consensus estimates. When participants were given a reason to attribute their increased ease of decision making to the valence of the frame and options, the effect disappeared (or even reversed, in the replication). Put another way, when participants are told their perceptions of decision ease may have been distorted, they no longer use them to inform their confidence and consensus estimates. Importantly, that the warning did not indicate the direction of the effect suggests a corrected attribution of ease, rather than simply demand, is driving the effect. However, one could say that, because the warning does tell participants that aspects of the decision could influence decision ease, participants may still perceive ease to be useful information (although our attenuation suggests they did not).

An older, alternative version of this paradigm (Schwarz, Bless, Strack, Klumpp, Rittenauer-Schatka, & Simons, 1991) avoids this point by having participants attribute their decision ease to something outside of the decision itself (e.g., music playing in the background). We chose the above approach (see Cesario, Grant, & Higgins, 2004), because that alternative unnecessarily deceives participants and appears statistically unreliable (e.g., most of the reported test statistics are not significant). Moreover, given either method could, in concept, preserve the diagnosticity of ease for participants (i.e., both "The valence of that decision must have really drawn my focus and helped me find my true preferences" and "That music must have really cleared my mind and helped me find my true preferences") yet do not seem to do so, we preferred to employ the more reliable version. (At the request of journal reviewers, we attempted to rerun Study 7 with this type of approach but, in line with the absence of evidence in the original paper, we did not find the predicted attenuation of our effect. Full results are available in the Appendix.)

Having found support for our proposed moderator of decision ease in Studies 6 and 7, we next turn to rule out alternatives to this explanation of our effect. Studies 8 and 9 explain and address the possibility of response substitution (Gal & Rucker, 2011) driving our results, and Study 10 attempts to dispel concerns that a simple language effect could be at fault.

Study 8

Perhaps participants in mismatched trials were answering a different question than the one asked, in order to convey a strong opinion they are holding. When shown a pair of faces, for example, participants may want to answer "how attractive are these faces?" but are instead asked only "how confident are you in your decision?" In order to convey to the experimenter that they see a clear difference between the attractive and unattractive pairs of faces, they substitute their response to the former when answering the latter. Doing so would generate results consistent with our predictions without the psychology behind it. Therefore, in Study 8, we modify our standard study design to give some participants a chance to express these opinions. If response substitution is driving our effects, then participants who have the opportunity to provide their attitudes towards the stimuli should show no matching effects on our usual follow-up questions.

Method

Participants. We decided in advance to recruit 400 participants from Amazon Mechanical Turk and received 401 participants (M_{age} = 34.4, 52.6% female).

Materials, procedure, design. From Study 1's materials, we selected the two pairs of attractive faces and the two pairs of unattractive faces that showed the largest attribute matching effect. All participants saw these four pairs and, to conserve statistical power, all participants were asked which woman they would choose (i.e., we did not manipulate decision frame in this study). All participants reported how easy the choice was to make and how confident they were in their choice. However, for half of participants, after choosing a woman but before reporting decision ease, they were asked to rate the attractiveness of the women in the pair (1 = very unattractive, 7 = very attractive). After completing all four trials, participants reported their age and gender, and completed a brief attention check. (Pre-registration and data: https://osf.io/zg6b5/.)

Results and Discussion

Sixty-eight participants (17.2%) failed the attention check and were excluded from analyses.

Ease. With the remaining participants, as in Study 1, we ran an OLS regression, clustering standard errors at the participant level. We defined two Level-1 variables (*attract*: -0.5 = no attractiveness question, 0.5 = attractiveness question; and *option*: -0.5 = unattractive, 0.5 = attractive), and included their interaction. Consistent with our hypothesis that our previous results were not driven by response substitution, we found only a significant effect of *option* (b = 1.13, SE = 0.10, p < .001). Whether participants were able to express their opinions about the women's attractiveness beforehand had no effect on its own (b = -0.06, SE = 0.11, p = .577) nor did it moderate the *option* effect (b = 0.03, SE = 0.20, p = .866).

Confidence. We then repeated this analysis for reported decision confidence. Again, only *option* emerged as a significant predictor (b = 1.68, SE = 0.12, p < .001); allowing some participants to express their opinions beforehand had no effects on its own (b = .21, SE = .14, p = .153) or as a moderator (b = 0.22, SE = 0.24, p = .357).

Study 9

The goal of Study 9 was to begin to address a concern that attribute matching is merely a language effect. That is, put simply, if a question is phrased strangely, people will be put off. By this account, mismatched questions could elicit lower levels of confidence (for example) because they imply enthusiastic endorsement (e.g., "Yes, I want tripe for dinner") that the decision maker may not actually feel. That reluctance is then expressed through our follow-up measures. If attribute matching were driven by this language effect, then presenting participants with mismatched decisions phrased more hesitantly should adequately capture their reluctance, which would then not appear in our follow-up measures and eliminate the effect. However, we predict that is not the case, and attribute matching will emerge regardless of how un/enthusiastically the question is phrased. Pre-registration and data: https://osf.io/r2khx/).

Method

Participants. We recruited 508 participants from Amazon Mechanical Turk. We had pre-registered to recruit 500 in total.

Materials, procedure, and design. Participants saw the same dinner party scenario as in Study 5, in which they saw a pair of appealing foods and a pair of unappealing foods and made either a choice or rejection. However, half of participants were randomly assigned to reluctance conditions. For these participants, the study's instructions included a few sentences acknowledging the potential difficulty of the task: (e.g., for the choose condition) "You may find both options appealing and wish you could choose both; or you may find both options unappealing and wish you could choose neither. We understand. Please, imagine you **had to** choose one, and tell us which you would choose if you were forced to choose one." Then, the two decisions were prefaced with similar language. Finally, when participants were reminded of their choice before reporting their confidence, it was phrased as "If you had to, you would reluctantly choose [reject]..." After completing both trials, participants completed the same attention check as in previous studies and reported whether they were vegetarian or on a diet.

Results

Sixty-eight (13.4%) participants failed the attention check and were excluded from analyses. With the remaining participants, and using trial as the unit of analysis, we regressed confidence ratings on *frame* (-0.5 = choose, 0.5 = reject), *option* (-0.5 = negative, 0.5 = positive), *reluctance* (-0.5 = no, 0.5 = yes), all possible two-way interactions, and the three-way *frame x option x reluctance* interaction. We clustered standard errors at the participant level. Surprisingly, we found no significant effects, with the exception of *frame* (b = 0.57, SE = 0.16, p < .001), including our standard *frame X option* interaction (b = 0.37, SE = 0.35, p = .295) or the (as predicted) three-way interaction (b = -0.61, SE = 0.70, p = .388).

However, looking to the simple effects, we do still see evidence of the reversing pattern from previous studies. In the standard phrasing conditions, participants on positive trials were more confident in their decision when choosing the better food (M = 7.71, SD = 1.44) than those rejecting the worse food (M = 6.80, SD = 2.17), $t(197) = 3.82$, $p < .001$, an effect that reversed for negative trials ($M_{choose} = 4.46$, $SD_{choose} = 2.75$ vs. $M_{reject} = 6.38$, $SD_{reject} = 2.45$; $t(197) = 5.70$, $p < .001$). Similarly, and critically counter to the prediction from a language-based account, the same pattern emerged when the decisions were phrased reluctantly: On positive trials, participants choosing the better food were more confident in their decision (M = 7.64, SD = 1.88) than participants rejecting the worse food (M = 7.03, SD = 1.98), $t(239) = 2.59$, $p = .010$, an effect that reversed for negative trials ($M_{choose} = 5.15$, $SD_{choose} = 2.79$ vs. $M_{reject} = 6.73$, $SD_{reject} = 2.35$; $t(239) = 5.06$, $p < .001$).

The absence of the reversal interaction we have seen in previous studies precludes a more definitive comment on the possibility of language effects driving attribute matching. However, that framing the decision reluctantly did not appear to meaningfully affect the pattern of means above, suggests that it is not likely a driving force. Nevertheless, in Study 10, we employ an alternative approach to answer this question of whether simple language effects could be behind this effect.

Study 10

So far, we have shown that matching positively (negatively)-framed decisions with positive (negative) outcomes. In Study 3, we demonstrated the role of attribute intensity, wherein less intensely positive or negative words showed a smaller (but still significant) matching effect. However, an alternative explanation for this finding is a simple language effect: Matching slightly positive words with fully positive decision frames (e.g. "I choose ____") may actually be a mismatch of its own. Perhaps if a less intense frame were matched with the less intense options, attribute matching would appear just as strongly. As in Study 9, this would suggest an alternative, less interesting explanation of our effect: People dislike oddly framed decisions. To dispel this concern, we adopt the multiple levels of option valence used in Study 3, but add additional levels of decision-frame to match. We predict that attribute matching will emerge as before, with no additional effect of frame intensity across levels of option valence. Pre-registration and data: https://osf.io/a25yr/).

Method

Participants. We recruited 504 participants from Amazon Mechanical Turk. We had pre-registered to recruit 500 in total.

Materials, procedure, and design. Participants saw four pairs of outdoor scenes that were pretested by 50 Amazon Mechanical Turk workers to be very appealing (e.g., beach vistas), slightly appealing, slightly unappealing, or very unappealing (e.g., landfills). Given the mixed results from our reluctance manipulation in the previous study, in Study 10, we manipulated decision frame intensity by randomly assigning participants to make one of four judgments: which image is more gorgeous, appealing, unappealing, or ugly. After each judgment, participants reported how easy the decision felt to make, as in Studies 2 and 6. After completing all four trials, participants completed an attention check: a multiple-choice question that asked what type of images they had viewed in the study.

Results

Three participants (0.6%) participants failed the attention check and were excluded from analyses. With the remaining participants, and using trial as the unit of analysis, we regressed confidence ratings on *frame* (-0.5 = positive, 0.5 = negative), *option* (-0.5 = negative, 0.5 = positive), *f_intensity* (-0.5 = low, 0.5 = high), *o_intensity* (-0.5 = low, 0.5 = high), all possible two-way and three-way interactions, and the four-way *frame x option x f_intensity x o_intensity* interaction. We clustered standard errors at the participant level.

Overall, we found a significant effect of *option* (b = 0.69, SE = 0.07, $p < .001$) with no effect of *frame* (b = 0.00, SE = 0.10, p = .996), as well as our predicted interaction between the two (b = -1.63, SE = 0.15, $p < .001$). As before, on trials with positive images, participants asked to make a positive judgment reported the decision was easier to make (M = 5.16, SD = 1.18) than participants making a negative judgment (M = 4.36, SD = 1.32), $t(499) = 7.16$, $p < .001$, an effect that reversed for trials with negative images ($M_{positive} = 3.70$, $SD_{positive} = 1.63$ vs. $M_{negative} = 4.48$, $SD_{negative} = 1.34$; $t(499) = 5.87$, $p < .001$).

In line with the results from Study 3, this *frame x option* interaction was also moderated by *frame_intensity* (b = -0.77, SE = 0.29, p = .009) and *option_intensity* (b = -1.48, SE = 0.28, $p < .001$), although the unexpectedly and consistently high ratings for the "good" pair of images (the second bar in each set) makes this result difficult to interpret. Crucially, though, in contrast to a language-based account, the four-way interaction was not significant (b = -0.65, SE = 0.56, p = .249). The full set of means is plotted in Figure 7, but an important pattern emerges: Although matching effects are smaller for less intense frames or less intense options (per Study 3), there is no language-based matching effect. Were that the case, the highest ease ratings would appear for gorgeous/great, appealing/good, unappealing/bad, and ugly/terrible, which did not obtain.

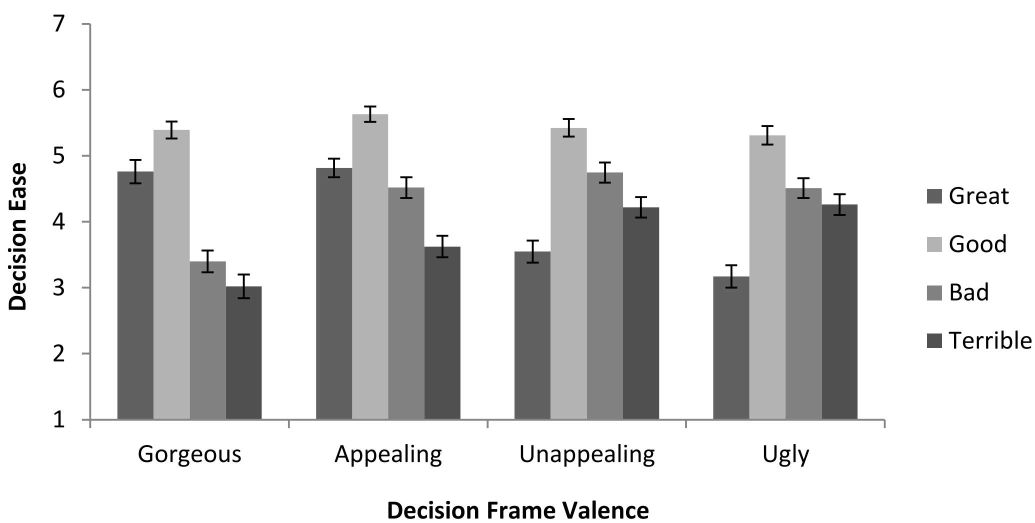

Figure 7. Matching the frame's valence intensity with that of the options does not enhance the effect in Study 10.

In addition to those investigated in the previous studies, another alternative explanation often put forth for matching effects is positive mood, rather than decision ease (e.g., Cesario, et al., 2004). Although we did not directly measure mood in our studies, most of our studies included explicitly negative stimuli (e.g., the word "murder" or liverwurst sausage), which do not typically induce a positive mood. Accordingly, this too seems like a relatively unlikely account for our findings. Therefore, having ruled several alternative explanations, we now turn to examining a more realistic context: consumer decision-making.

Study 11

Up until this point, our focus has been on demonstrating the existence of this attribute matching effect and ruling in or out possible explanations. However, a reasonable criticism thus far would be that our stimuli have been rather abstract (e.g., pairs of good and bad words). Additionally, one could argue that people are frequently faced with choices, but far less often face rejections to make. With the following study, we aim to begin addressing these concerns by demonstrating attribute matching in a realistic decision context with realistic outcomes.

We do this by turning to the online marketplace, specifically for electronics. Because products are frequently displayed in a list form online, with minimal information, and electronic goods (e.g., computers, smartphones, cameras) vary on so many dimensions, some online retailers (e.g., Best Buy, NewEgg) offer a comparison feature for their site's visitors. This feature allows visitors to isolate a small number of options they are considering and view them side-by-side on a separate comparison page. In Study 11, we vary how the next step in the purchase process is phrased, as a more ecologically valid substitute for our previous decision frame manipulations. Despite this change of context, we predict that attribute matching will still emerge. Pre-registration and data: https://osf.io/s9wp2/).

Method.

Participants. We pre-registered to recruit and successfully recruited 400 participants from Amazon Mechanical Turk.

Materials, procedure, and design. Participants were asked to imagine they were looking to buy a new phone online, had narrowed down their options to two phones, and had put them both in their cart so that they could better compare the phones' features. They were then shown a pair of (at the time) new iPhones (positive outcomes) and a pair of flip phones (negative outcomes) in a randomized order, receiving the scenario description at the start of each pair. For the decision, some participants were asked "Which phone would you KEEP in your cart to purchase?" (positive frame); others were asked "Which phone you would DELETE from your cart before proceeding to purchase?" (negative frame). After each decision, participants reported their certainty in their selection on a 9-point scale as in previous studies. After making both decisions, participants proceeded to make four other decisions in unrelated domains. Finally, participants completed the attention check from previous studies and reported their highest level of completed education.

Results and Discussion

Fifty-nine participants (14.8%) of participants failed the attention check and were excluded from analyses. As in previous studies, using trial as the unit of analysis, we regressed reported certainty on *frame* (-0.5 = choose, 0.5 = reject), *option* (-0.5 = negative, 0.5 = positive), and the interaction between them. We clustered standard errors at the participant level. The main effect of f*rame* was not significant (b = -0.11, SE = 0.16, p = .482); that of *option* was (b = 0.28, SE = 0.11, p = .010). The interaction was significant (b = -0.89, SE = 0.21, p < .001), but not fully with the predicted pattern: On positive trials, participants choosing the better word reported greater certainty (M = 7.59, SD = 1.48) than participants rejecting the worse word (M = 7.04, SD = 1.78), $t(339) = 3.12$, $p = .002$. Unlike previous studies, however, this effect only marginally reversed for negative trials ($M_{choose} = 6.87$, $SD_{choose} = 1.83$ vs. $M_{reject} = 7.21$, $SD_{reject} = 1.86$; $t(339) = 1.66$, $p = .097$).

Although we did not perfectly replicate our previous crossover interaction, the pattern still persisted, even in this consumer domain with different decision frames. One reason we may have found weaker effects for the negative trial is that having a flip phone, while not the preferred outcome for most, is not actually aversive. Drawing on the results from Study 3, which showed smaller effects for slightly negative options compared to very negative options, perhaps we should have anticipated this outcome.

This then raises an interesting problem for applying attribute matching to consumer decision making: By definition, goods that persist in the market have some positive utility to them. That is, one could say that there are no negative goods (to match or mismatch with differently valenced frames). We would disagree. First, Study 2 demonstrated that this effect need not be limited to valence matching. Second, and more important, goods can have utility to them without being desired. For example, in their work on disutility from lose-lose decisions, Botti, Orfali, and Iyengar (2009) looked at parents whose newborn children had life-threatening complications. The various treatment options—a risky procedure, wait and see, remove life support—all certainly offered potential benefits to the worried parents and their sick children, but all parties would be much happier if they were not needed in the first place. One can find analogs in everyday life (e.g., healthy foods or different exercise routines when reluctantly on a diet, administrative or busywork tasks, unpleasant medical procedures), hence we believe that Study 11 was underpowered to detect this smaller difference, rather than concede that attribute matching cannot appear in real-world contexts.

General Discussion

No one wants to face undesirable options, but, given that preferences are malleable, the present paper demonstrates that people might still feel good about their decisions amongst them. We propose that this can happen through attribute matching: When a salient attribute of the decision frame matches a salient attribute of the options, the decision feels easier to make, which increases reported confidence and perceived consensus. In Study 1, we showed that participants were more confident in their preference and perceived greater consensus around it when they were choosing from attractive faces or rejecting from unattractive faces than when the frame and the options did not match in valence. In Study 2, we extended this finding to objective judgments. Despite experiential information now having to compete with much more declarative information, we still found evidence of attribute matching with calorie judgments in Study 2. Study 3 found that this matching effect emerged even with only slightly valenced options. Studies 4-5 generalized our effects on confidence and consensus to a variety of other attitude types. Study 6 directly tested a possible mechanism by measuring decision ease and response time and finding strong evidence of attribute matching and mediation of confidence estimates; hence, in Study 7, notifying participants that the valence of the frame and options may affect their decision ease appeared to correct their misattribution of that ease and eliminate the effect. Studies 8-10 ruled out alternative explanations, and Study 11 examined this effect in a more realistic context.

A fruitful avenue for building off this work could come in the form of additional mechanisms. That is, although we show confidence estimates drive perceived consensus, and decision ease drives confidence estimates, the question remains as to what drives decision ease. Given the subjective nature of most of our stimuli, cognitive dissonance (Festinger, 1957) decreasing on matched trials and increasing on mismatched trials is certainly involved in some form: One feels uncomfortable in rejecting a perfectly good option and better being able to choose one. However, this account cannot accommodate our results from Study 2, in the objective domain: stating that burgers have more calories than pizza should be affectively equivalent to stating pizza has fewer calories than burgers, yet we show strong effects on decision ease in both cases (which is strongly related to decision conflict or dissonance), hence further research may prove useful.

Another possibility is that the decision frame may cue participants to selectively recall, search, and process matching information (Shafir, 1993; see Nickerson (1998) for a review on a related topic, confirmation bias). On mismatched trials, however, there is less of this desired information available, increasing decision difficulty. A third option could involve cognitive switching costs (Pecher, Zeelenberg, & Barsalou, 2003; see also Messner &Vosgerau, 2010; Allport & Wylie, 1999). That is, being given a positive frame may facilitate subsequent processing of positive options (a match) and hinder processing of negative options (a mismatch), which manifests as increased or decreased decision ease.

There are other existing literatures that, despite considering themselves distinct from each other, may be, in fact, quite similar to both each other and to the work presented above. However, none of them answers the key questions posed here. Therefore, another important stream of future research could look into connecting of some of these otherwise disparate findings in the literature to ours. For example, work on regulatory fit (Higgins, 2000) claims that "when [people] use goal pursuit means that fit their regulatory orientation" (p. 1219) this fit generates its own utility (Higgins, 2003), increasing decision confidence (Cesario, Grant, & Higgins, 2004) and satisfaction (Idson, Liberman, & Higgins, 2000). Construal level theory (Trope & Liberman, 2000) research has shown when an option's features (construals) match each other, decision makers view those options more positively (Fujita et al, 2006; Kim, Rao, & Lee, 2009; Todorov, Goren, & Trope, 2007). Beyond these two literatures, a number of papers in other domains appear to show a similar pattern: e.g., rounded, easy to calculate numbers facilitate affect-driven decisions (Wadhwa & Zhang, 2015) and option information in the same units as the decision is utilized more in decision making (Nowlis & Simonson, 1997; Slovic, Griffin, & Tversky, 1990;). It could be that, in abstract, these papers and literatures are each their own form of attribute matching: each matching attributes from different parts of the decision process, with positive outcomes for the decision maker. If any mechanism is proposed and tested empirically in the above work, it is decision ease. Hence the case of attribute matching we present here is a pure, robust, reliable form of what could be at the core of these important literatures, and better understanding the mechanism behind it would certainly be valuable.

A third, more downstream, focus of future research could include greater investigation of the consequences of these changes in decision ease. The attitude certainty literature suggests several important behavioral outcomes, such as resistance to change (Petrocelli, Tormala, & Rucker, 2007), persistence over time (Bassili, 1996), and likelihood of acting on that attitude (Tormala & Petty, 2002) when the attitude is more certain, on matched trials. The regulatory fit literature suggests that positive decision ease on matched trials may spillover into evaluating the selected option more positively. There are clearly a number of exciting, important investigations to be built from this initial finding.

References

Allport, A. & Wylie, G. (1999). Task-switching: Positive and negative priming of task set. J. Duncan, A. Treisman , & E. Edmunds (Eds), Attention, Space, and Action: Studies in Cognitive Neuroscience. New York, Oxford University Press.

Alter, A. L., & Oppenheimer, D. M. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychology Review,20,* 1414-1420.

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*, 637–653.

Beattie, J., Baron, J., Hershey, J. C., & Spranca, M. D. (1994). Psychological determinants of decision attitude. *Journal of Behavioral Decision Making*, *7*, 129–144.

Bellezza, F. S., Greenwald, A. G., & Banaji, M. R. (1986). Words high and low in pleasantness as rated by male and female college students. *Behavior Research Methods, Instruments, & Computers*, *18*, 299-303.

Bizer, G. Y., & Petty, R. E. (2005). How we conceptualize our attitudes matters: The effects of valence framing on the resistance of political attitudes. *Political Psychology*, *26*, 553-568.

Botti, S., & Iyengar, S. S. (2004). The psychological pleasure and pain of choosing: when people prefer choosing at the cost of subsequent outcome satisfaction. *Journal of Personality and Social Psychology*, *87*, 312-326.

Botti, S., & Iyengar, S. S. (2006). The dark side of choice: When choice impairs social welfare. *Journal of Public Policy & Marketing*, *25*, 24-38.

Botti, S., & McGill, A. L. (2006). When choosing is not deciding: The effect of perceived responsibility on satisfaction. *Journal of Consumer Research*, *33*, 211-219.

Botti, S., Orfali, K., & Iyengar, S. S. (2009). Tragic choices: Autonomy and emotional responses to medical decisions. *Journal of Consumer Research*, *36*, 337-352.

Burger, J. (1989). Negative reactions to increases in perceived personal control. *Journal of Personality and Social Psychology*, *56*, 246–256.

Burson, K. A., Larrick, R. P., & Lynch, J. G. (2009). Six of one, half dozen of the other expanding and contracting numerical dimensions produces preference reversals. *Psychological Science*, *20*, 1074-1078.

Campbell, T., O'Brien, E., Van Boven, L., Schwarz, N., & Ubel, P. (2014). Too much experience: a desensitization bias in emotional perspective taking. *Journal of Personality and Social Psychology*, *106*, 272-285.

Cesario, J., Grant, H., & Higgins, E. T. (2004). Regulatory fit and persuasion: Transfer from "feeling right." *Journal of Personality and Social Psychology, 86*, 388-404.

Dawes, R.M. (1990). The potential nonfalsity of the false consensus effect. R.M. Hogarth (Ed.), Insights in Decision Making: A Tribute to Hillel J. Einhorn, University of Chicago Press, Chicago.

Dhar, R. (1997). Context and task effects on choice deferral. *Marketing Letters*, *8*, 119-130.

Egan, L. C., Santos, L. R., & Bloom, P. (2007). The origins of cognitive dissonance evidence from children and monkeys. *Psychological Science*, *18*, 978-983.

Evans, A. M., Dillon, K. D., Rand, D.G. (2015). Fast but not intuitive, slow but not reflective: Decision conflict drives reaction times in social dilemmas. *Journal of Experimental Psychology: General*, *144*, 951-966.

Fazio, R. H., & Zanna, M. P. (1978). Attitudinal qualities relating to the strength of the attitude-behavior relationship. *Journal of Experimental Social Psychology*, *14*, 398-408.

Fujita, K., Henderson, M., Eng, J., Trope, Y., & Liberman, N. (2006). Spatial distance and mental construal of social events. *Psychological Science, 17*, 278–282.

Gal, D., & Rucker, D. D. (2011). Answering the unasked question: Response substitution in consumer surveys. *Journal of Marketing Research*, *48*, 185-195.

Gilovich, T., & Medvec, V. H. (1995). The experience of regret: what, when, and why. *Psychological Review*, *102*, 379-395.

Goldstein, N. J., Cialdini, R. B., & Griskevicius, V. (2008). A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of Consumer Research*, *35*, 472-482.

Gross, S., Holtz, R., & Miller, N. (1995). *Attitude certainty*. In R. E. Petty & J. A. Krosnick (Eds.), *Attitude Strength: Antecedents and Consequences* (pp. *215–245*). Mahwah, NJ: Lawrence Erlbaum Associates.

Higgins, E. T. (2000). Making a good decision: value from fit. *American Psychologist*, *55*, 1217-1230.

Higgins, E. T. (2005). Value from regulatory fit. *Current Directions in Psychological Science*, *14*, 209-213.

Horcajo, J., Petty, R. E., & Brinol, P. (2010). The effects of majority versus minority source status on persuasion: a self-validation analysis. *Journal of Personality and Social Psychology*, *99*, 498-512.

Houston, D. A., Sherman, S. J., & Baker, S. M. (1991). Feature matching, unique features, and the dynamics of the choice process: Predecision conflict and postdecision satisfaction. *Journal of Experimental Social Psychology, 27*, 411-430.

Idson, L. C., Liberman, N., & Higgins, E. T. (2000). Distinguishing gains from nonlosses and losses from nongains: A regulatory focus perspective on hedonic intensity. *Journal of Experimental Social Psychology*, *36*, 252-274.

Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases.* New York: Cambridge University Press.

Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. *American Psychologist*, *39*, 341-350.

Kim, H., Rao, A. R., & Lee, A. Y. (2009). It's time to vote: The effect of matching message orientation and temporal frame on political persuasion. *Journal of Consumer Research, 35*, 877-889.

Krueger, J., & Clement, R. W. (1994). The truly false consensus effect: an ineradicable and egocentric bias in social perception. *Journal of Personality and Social Psychology*, *67*, 596-610.

McGlone, M. S., & Tofighbakhsh, J. (2000). Birds of a feather flock conjointly (?): Rhyme as reason in aphorisms. *Psychological Science, 11*, 424-428.

Meloy, M. G., & Russo, J. E. (2004). Binary choice under instructions to select versus reject. *Organizational Behavior and Human Decision Processes, 93,* 114-128.

Messner, C., & Vosgerau, J. (2010). Cognitive inertia and the implicit association test. *Journal of Marketing Research*, *47*, 374-386.

Nagpal, A., & Krishnamurthy, P. (2008). Attribute conflict in consumer decision making: The role of task compatibility. *Journal of Consumer Research*, *34*, 696-705.

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, *2*, 175-220.

Nowlis, S. M., & Simonson, I. (1997). Attribute-task compatibility as a determinant of consumer preference reversals. *Journal of Marketing Research*, 205-218.

Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, *45*, 867-872.

Park, C. W., Jun, S. Y., & MacInnis, D. J. (2000). Choosing what I want versus rejecting what I do not want: An application of decision framing to product option choice decisions. *Journal of Marketing Research*, *37*, 187-202.

Pecher, D., Zeelenberg, R., & Barsalou, L. W. (2003). Verifying different-modality properties for concepts produces switching costs. *Psychological Science*, *14*, 119-124.

Perlmuter, L. C., & Monty, R. A. (1977). The importance of perceived control: Fact or fantasy? Experiments with both humans and animals indicate that the mere illusion of control significantly improves performance in a variety of situations. *American Scientist*, 759-765.

Petrocelli, J. V., Tormala, Z. L., & Rucker, D. D. (2007). Unpacking attitude certainty: attitude clarity and attitude correctness. *Journal of Personality and Social Psychology*, *92*, 30-41.

Petty, R. E., & Krosnick, J. A. (1995). Attitude strength: Antecedents and consequences. Mahwah, NJ: Erlbaum.

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.

Pochon, J. B., Riis, J., Sanfey, A. G., Nystrom, L. E., & Cohen, J. D. (2008). Functional imaging of decision conflict. *The Journal of Neuroscience*, *28*, 3468-3473.

Reber, R., Winkielman, P., & Schwarz, N. (1998). Effects of perceptual fluency on affective judgments. *Psychological Science*, *9*, 45-48.

Ross, L., Greene, D., & House, P. (1977). The “false consensus effect”: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology*, *13*, 279-301.

Rucker, D. D., Tormala, Z. L., Petty, R. E., & Briñol, P. (2014). Consumer conviction and commitment: An appraisal-based framework for attitude certainty. *Journal of Consumer Psychology*, *1*, 119-136.

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

Schwarz, N. (1998). Accessible content and accessibility experiences: The interplay of declarative and experiential information in judgment. *Personality and Social Psychology Review*, *2*, 87-99.

Shafir, E. (1993). Choosing versus rejecting: Why some options are both better and worse than others. *Memory and Cognition, 21*, 546-556.

Shafir, E., Simonson, I., & Tversky, A. (1993). Reason-based choice. *Cognition*, *49*, 11-36.

Simmons, J. P., & Nelson, L. D. (2006). Intuitive confidence: choosing between intuitive and nonintuitive alternatives. *Journal of Experimental Psychology: General*, *135*, 409-428.

Simmons, J. P., Nelson, L. D., and Simonsohn, U. (2012). A 21 word solution. *Dialogue: The Official Newsletter of the Society for Personality and Social Psychology, 26*, 4-7.

Slovic, P., Griffin, D., & Tversky, A. (1990). *Compatibility effects in judgment and choice*. In H. J. Einhorn & R. M. Hogarth (Eds.), *Insights in decision making: A tribute to Hillel J. Einhorn*, (pp. 5-27). Chicago, IL: University of Chicago Press.

Sokolova, T., & Krishna, A. (in press). Take It or Leave It: How Choosing versus Rejecting Alternatives Affects Information Processing. *Journal of Consumer Research*.

Stotland, E., & Blumenthal, A. L. (1964). The reduction of anxiety as a result of the expectation of making a choice. *Canadian Journal of Psychology, 18*, 139-145.

Taylor, S. E., Lichtman, R. R., & Wood, J. V. (1984). Attributions, beliefs about control, and adjustment to breast cancer. *Journal of Personality and Social Psychology*, *46*, 489-502.

Thaler, R. H., & Sunstein, C. R. (2008). Nudge: Improving Decisions about Health, Wealth, and Happiness. New Haven, CT: Yale University Press.

Todorov, A., Goren, A., & Trope, Y. (2007). Probability as a psychological distance: Construal and preferences. *Journal of Experimental Social Psychology, 43*, 473-482.

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

Tormala, Z.L., Petty, R.E., & 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. (2007). Attitude certainty: A review of past findings and emerging perspectives. *Social and Personality Psychology Compass, 1,* 469-492.

Trope, Y., & Liberman, N. (2000). Temporal construal and time-dependent changes in preference. *Journal of Personality and Social Psychology, 79*, 876-889.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*, 453-458.

Tversky, A., & Shafir, E. (1992). The disjunction effect in choice under uncertainty. *Psychological science*, *3,* 305-309.

Van Boven, L., & Loewenstein, G. (2003). Social projection of transient drive states. *Personality and Social Psychology Bulletin*, *29*, 1159-1168.

Visser, P. S., & Mirabile, R. R. (2004). Attitudes in the social context: The impact of social network composition on individual-level attitude strength. *Journal of Personality and Social Psychology*, *87*, 779–795.

Wadhwa, M., & Zhang, K. (2015). This number just feels right: The impact of roundedness of price numbers on product evaluations. *Journal of Consumer Research*, *41*, 1172-1185.

Appendix

First, we present the decisions made by participants in Study 2, split by condition, to show our manipulation did not significantly influence what participants selected. Next, we present the results of seven studies that were not included in the final paper: The first was excluded out of length and relevance concerns and is referred to in the body of the paper. The second study was successful, but excluded out of length and redundancy concerns. The next two were ambiguously successful and the final two unsuccessful, but the set may still be informative.

1. Trial-Level Decisions from Study 2

For ease of comparison to those of the choose condition, the reject conditions' selections have been reversed to show their de facto choice. Note that on only one trial (trial nine) of twenty did the majority preference differ between the two conditions (likely due to being so close to 50%).

Trial Number (valence)	**Frame**	**% Choosing A**	**% Choosing B**
1 (extremely negative)	Choose	41.3	58.7
	Reject	42.4	57.6
2 (extremely negative)	Choose	68.0	32.0
	Reject	85.4	14.6
3 (extremely negative)	Choose	62.7	37.3
	Reject	64.9	35.1
4 (extremely negative)	Choose	38.7	61.3
	Reject	35.1	64.9
5 (extremely negative)	Choose	70.7	29.3
	Reject	70.2	29.8
6 (slightly negative)	Choose	72.7	27.3
	Reject	66.2	33.8
7 (slightly negative)	Choose	47.3	53.7
	Reject	28.5	71.5
8 (slightly negative)	Choose	42.0	58.0
	Reject	41.1	58.9
9 (slightly negative)	Choose	50.7	49.3
	Reject	41.1	58.9
10 (slightly negative)	Choose	33.3	66.7
	Reject	23.2	76.8
11 (slightly positive)	Choose	36.7	63.3
	Reject	22.5	77.5
12 (slightly positive)	Choose	44.0	56.0
	Reject	22.5	77.5
13 (slightly positive)	Choose	76.7	23.3
	Reject	69.5	30.5
14 (slightly positive)	Choose	64.7	35.3
	Reject	82.8	17.2
15 (slightly positive)	Choose	56.0	44.0

	Reject	61.6	38.4
16 (extremely positive)	Choose	63.3	36.7
	Reject	66.2	33.8
17 (extremely positive)	Choose	71.3	28.7
	Reject	62.3	37.7
18 (extremely positive)	Choose	48.7	51.3
	Reject	39.1	60.9
19 (extremely positive)	Choose	75.3	24.7
	Reject	67.5	32.5
20 (extremely positive)	Choose	66.0	34.0
	Reject	84.8	15.2

2.1. Schwarz, et al (1991) Misattribution Study

Method

Participants. We decided in advance to recruit 500 participants from Amazon Mechanical Turk and received 513 participants (M_{age} = 30.6, 60.2% female).

Materials, procedure, design. This study was identical to Study 4, with two exceptions: First, all text was shown on a teal background. The colored background was in place of the background music used by Schwarz, et al (1991), since providing background music would be difficult to implement reliably online. Second, instead of being told about the role of frame and option valence, participants in the warned condition were told that "past research suggests that making decisions on teal backgrounds can change how easy that decision feels to make." (Pre-registration and data: https://osf.io/h7sqb/.)

Results.

Sixty-seven participants (13.1%) failed the attention check and were excluded from analyses.

Confidence. First, for confidence, we fit the same model as in Study 4 but with *warning* (-0.5 = unwarned, 0.5 = warned) as a third predictor and its two-way interactions with *option* and *frame* and the three-way interaction. *Frame* (b = 0.54, SE = 0.10, p < .001) and *option* (b = 0.54, SE = 0.06, p < .001) were both significant predictors of reported confidence. As in previous studies (including Study 4), these effects were qualified by the predicted crossover interaction (b = -3.29, SE = 0.12, p < .001). However, unlike Study 4, the predicted three-way interaction was not significant, b = -0.13, SE = 0.22, p = .567. For completeness, we discuss the unwarned and warned conditions separately, below.

In the unwarned conditions, without the warning manipulation, attribute matching replicated as predicted: on positive trials, participants reported greater confidence when choosing (M = 7.17) than rejecting (M = 6.12), $t(233)$ = 6.32, p < .001; whereas on negative trials, participants reported greater confidence when rejecting (M = 7.17) than choosing (M = 4.98), $t(223)$ = 13.94, p < .001. In the warned conditions, the results were quite similar: for positive trials participants reported greater confidence when they were choosing (M = 7.43) than rejecting (M = 6.26), $t(219)$ = 7.38, p < .001; whereas on negative trials, participants reported greater confidence when rejecting (M = 7.44) than choosing (M = 5.25), $t(219)$ = 14.94, p < .001.

Consensus. We then fit the same model for consensus estimates. Again, both f*rame* (b = 3.33, SE = 0.88, p < .001) and option valence (b = 1.74, SE = 0.62, p = .006) significantly predicted estimated consensus, as well as, our predicted crossover interaction was highly significant (b = -2.73, SE = 0.61, p < .001). However, again, the predicted three-way interaction was not significant, *b* = -3.49, SE = 2.49, $p = .161$. For completeness, we again discuss the unwarned and warned conditions separately, below.

In the unwarned conditions, without the warning, attribute matching replicated as predicted: on positive trials, participants estimated higher consensus when choosing (M = 64.1%) than rejecting (M = 55.1%), $t(223) = 6.20$, $p < .001$; whereas on negative trials, participants estimated higher consensus when rejecting (M =66.0%) than choosing (M = 49.5%), $t(223) = 11.22$, $p < .001$. In the warned conditions, the results were again quite similar: on positive trials, participants estimated higher consensus when choosing (M = 67.4%) than rejecting (M = 55.8%), $t(219) = 7.45$, $p < .001$; whereas on negative trials, participants estimated higher consensus when rejecting (M =68.7%) than choosing (M = 51.2%), $t(219) = 10.91$, $p < .001$.

2.3. Punishments Study

We also ran a study using purely negative stimuli. Although Study 2-4 also have negative stimuli (negative words), we wanted to verify that attribute matching would still hold for negative stimuli outside of the artificial domain of those studies (i.e., word preferences). In this study, we presented participants with pairs of very aversive (e.g., cleaning public bathrooms) and less aversive (e.g., doing your laundry) tasks in either the “choose” or “reject” frame used in Studies 1-4. (Pre-registration and data: https://osf.io/bg4j6/.)

Method

Participants. We recruited 1,152 participants from the same private research company, and as part of the same larger project, as Study 1. We determined this sample size to be half of all participants in a survey session that required 1,500 attention check passes in a separate survey. We did not analyze the data until that threshold had been met and data collection ceased.

Materials, procedure, design. Participants viewed three pairs of very unpleasant tasks and three pairs of slightly unpleasant tasks, in a randomized order. The tasks’ unpleasantness was determined by a pretest, in which 52 Amazon Mechanical Turk participants were asked to rate each task’s unpleasantness separately. For these six pairs, participants were randomly assigned to report which task they would rather do or which task they would rather not do. After making each selection, participants completed the same confidence and consensus measures from the previous studies. At the conclusion of the study, participants received a brief attention check and reported their age and gender.

Results

We first excluded 398 participants (34.5%) for failing the attention check.

Confidence. With the remaining participants, we ran an OLS regression, clustering standard errors at the participant level, with *frame*, for the valence of the decision frame (-0.5 = choose, 0.5 = reject), and *option,* for the valence of the decision’s options (0.5 = less unpleasant, -0.5 = more unpleasant), as well as their interaction. A main effect of *frame* was not significant ($b = -0.05$, SE = 0.08, $p = .533$), although a main effect of *option* was ($b = 0.40$, SE = 0.05, $p < .001$). As predicted, however, our predicted interaction was highly significant ($b = -.62$, SE = 0.09, $p < .001$): for the less pleasant tasks, participants who were choosing were more confident in their decision (M = 7.93) than were the participants rejecting (M = 7.56), $t(752) = 11.09$, $p < .001$. For the more unpleasant tasks, the effect reversed ($M_{choose} = 7.22$ vs. $M_{reject} = 7.48$), $t(752) = 4.17$, $p < .001$).

Consensus. We then fit the same model for consensus estimates. As before, there was no main effect of *frame* ($b = -0.99$, SE = 087, $p = .252$), whereas there was for *option* ($b = -1.33$, SE = 0.33, $p = .033$). The predicted crossover interaction was only marginally significant ($b = -2.41$, SE = 1.25, $p = .054$): for less pleasant tasks, participants who were choosing the better task estimated higher consensus around their decision (M = 63.8%) than were the participants rejecting the worse task (M = 61.6%), $t(752) = 2.07$, $p = .033$. For the more unpleasant tasks, however, there was no difference (M_{choose} = 63.9% vs. M_{reject} = 64.1%), $t(752) = 0.20$, $p = .844$).

2.4. First Calories Study

Study 2.4 was our first attempt at what is now Study 2. Participants saw pairs of high- and low-calorie foods and were either asked which food has more calories or which food has fewer, a question with a known answer. Deciding which high-calorie food has more calories and which low-calorie food has fewer were the matched trials. (Pre-registration and data: https://osf.io/ct3hm/.)

Table S1. Food pairs used in Study 2.4.

High-Calorie Foods	Low-Calorie Foods
Double cheeseburger vs. Medium pepperoni pizza	Cup of black coffee vs. Cup of black tea
Slice of cheesecake vs. Chocolate milkshake	Apple vs. Peach
Frosted cinnamon bun vs. Beef burrito	Banana vs. Glass of skim milk
3 fried chicken pieces vs. 12 oz. steak	Cup of plain yogurt vs. Glass of lemonade
12 buffalo chicken wings vs. Small cheese pizza	Orange vs. Slice of whole wheat bread

Method

Participants. We recruited 406 participants from Amazon Mechanical Turk (M_{age} = 29.1, 49.9% female). We had pre-determined to recruit 400 participants in total. We did not analyze the data until that threshold had been met and data collection ceased.

Materials, procedure, and design. Participants viewed five pairs of high-calorie foods and five pairs of low-calorie foods, in a randomized order. Foods were determined to be high- or low-calorie based on a pretest, in which 61 Amazon Mechanical Turk participants were asked to estimate the caloric content of each food separately. See Table 1 for all 10 food pairs. The pairs were then constructed such that the median estimates were approximately equal. For these 10 food pairs, study participants were randomly assigned to determine which food has more calories or which food has fewer. After making each selection, participants completed the same confidence and consensus measures from the previous studies. At the conclusion of the study, participants received a brief attention check ("Which food has more [fewer] calories? 3 grapes vs. Grilled cheese sandwich") and provided their age and gender.

Results and Discussion

Four participants (1.0%) failed the attention check and were excluded from analyses.

Confidence. With the remaining participants, we ran an OLS regression, clustering standard errors at the participant level, with *frame*, for the valence of the decision frame (-0.5 = more, 0.5 = less), and *option*, for the valence of the decision's options (0.5 = high-calorie foods, -0.5 = low-calorie foods), as well as their interaction. A main effect of *frame* ($b = 0.45$, SE = 0.10, $p < .001$) and *option* ($b = 0.16$, SE = 0.06, $p = .003$) emerged. More importantly, our predicted interaction was significant, as well ($b = 0.95$, SE = 0.11, $p < .001$): on high-calorie trials, participants who were asked which food has more calories gave higher confidence estimates (M = 6.51) than did participants who were asked which has fewer (M = 5.58), $t(400) = 8.07$, $p < .001$. For low-calorie trials, however, there was no difference between the "fewer" and "more" frames (M_{more} = 5.86, M_{fewer} = 5.89; $t(400) = 0.18$, $p = .856$).

Consensus. We then fit the same model for consensus estimates. Neither *frame* (b = 1.55, SE = 0.87, p = .073) and *option* (b = 0.66, SE = 0.54, p = .225) were significant predictors of consensus estimates; however, as before, the predicted crossover interaction was highly significant (b = 9.09, SE = 1.08, p < .001): participants who were asked which food had more calories gave higher consensus estimates for choices made for high-calorie food pairs (M = 62.6%) than for low-calorie pairs (M = 60.0%), $t(400)$ = 2.73, p = .007). However, this effect reversed for participants who were asked which foods had fewer calories (M_{more} = 59.7%, M_{fewer} = 62.6%; $t(400)$ = 2.93, p = .004). In addition, when we entered confidence measures into the model, the interaction effect for consensus estimates was reduced (from b = 9.09 to b = 4.17), z = 4.19, p < .001.

Study 2.4 shows that the effects of attribute matching are not limited to judgments of subjective preference, but emerge even when people are judging stimuli on objective dimensions. However, we did not consistently obtain our predicted crossover interaction. We therefore modified our stimuli to stimuli that were clearly, inherently, and intuitively high- or low-calorie in Study 5 and obtained the predicted crossover interaction. Before doing so however, we ran the following study.

2.5. Equating Low- and High-Calorie

Study 2.5 was almost identical to Study 2.4, except the low-calorie food pairs were increased in quantity (e.g., 10 apples; see Table 2) to be calorically equivalent to the high-calorie food pairs. In Study 2.4, the low-calorie food pairs were low on both attribute intensity and attribute salience. By increasing the volume of the low-calorie foods, we artificially increase their attribute intensity, while maintaining their low attribute salience. We therefore predicted that, despite the increase in actual calories, participants would still intuitively view the high-quantity/low-calorie pairs as simply low-calorie, and hence show similar results as in previous studies. (Pre-registration and data: https://osf.io/7krb6/.)

Table S2. Food pairs used in Study 2.4.

High-Calorie Foods	**Low-Calorie Foods**
Double cheeseburger vs. Medium pepperoni pizza	4 cups of plain yogurt vs. 4 glasses of lemonade
Slice of cheesecake vs. Chocolate milkshake	8 oranges vs. 8 slices of whole wheat bread
3 fried chicken pieces vs. 12 oz. steak	10 apples vs. 10 peaches

Method

Participants. We recruited 1,140 participants from the same private research company, and as part of the same larger project, as Study 1. We determined this sample size to be half of all participants in a survey session that required 1,500 attention check passes in a separate survey. We did not analyze the data until that threshold had been met and data collection ceased.

Materials, procedure, and design. Participants viewed three pairs of high-calorie foods and three pairs of low-calorie foods in high quantities, in a randomized order. The high-calorie food pairs were the same as from Study 2.4. The low-calorie/high-quantity pairs were the same food pairs as the low-calorie pairs from Study 2.4, but with the addition of a multiplier such that the perceived caloric content of each low-calorie/high-quantity pair matched that of a high-calorie pair (e.g., if the average perceived calorie content of a high-calorie pair was 500 calories, a low-calorie pair with an average of 50 calories would be multiplied by 10). Note that the multiplier was always the same within pairs (e.g., 10 apples vs. 10 peaches), to prevent participants from simply relying on quantity to make their judgments. As in Study 2.4, for these 6 food pairs, study participants were randomly assigned to determine which food has more calories or which food has fewer. After making each selection, participants completed the same confidence and consensus measures from the previous studies. Finally, participants received the same attention check as Study 2.4 at the conclusion of the study.

Results and Discussion

Forty-five participants (3.9%) failed the attention check and were excluded from analyses. Although these participants came from the same pool as Study 1, we attribute this much lower failure rate to our simpler attention check.

Confidence. We first fit the same mixed model as in Study 2.4 for confidence estimates. A main effect of *frame* emerged (b = 0.19, SE = 0.08 p = .015), although this did not occur for *option* (b = 0.04, SE = 0.03, p = .274). Crucially, however, the predicted interaction was significant (b =0.58, SE = 0.07 p < .001): on high-calorie trials, participants who were asked which food has more calories gave higher confidence estimates (M = 6.48) than did participants who were asked which has fewer (M = 6.00), t(1093) = 5.75, p < .001. As in Study 2.2, this difference was eliminated for low-calorie/high-quantity trials (M_{more} = 6.15, M_{fewer} = 6.26; t(1093) = 1.18, p = .237).

Consensus. We then fit the same model for consensus estimates. Again, a main effect of *frame* emerged (b = 2.23, SE = 0.76, p = .003) and but not for *option* (b = 0.28, SE = 0.38, p = .466); however, the predicted interaction was again significant (b = 5.84, SE = 0.76, p < .001): participants who were asked which food has more calories gave higher consensus estimates (M = 67.2%) than did participants who were asked which has fewer (M = 62.1%), t(1093) = 6.25, p < .001). However, similar to the confidence measure, this effect was eliminated for low-calorie/high-quantity trials (M_{more} = 64.0%, M_{fewer} = 64.7%; t(1093) = 0.79, p = .429). In addition, when we entered confidence measures into the model, the interaction effect for consensus estimates was reduced (from b = 5.84 to b = 2.42), z = 3.31, p < .001.

Study 2.5's results exhibit a strong likeness to Study 2.3's, despite all food pairs being high in calories. This result could be interpreted as participants still initially, intuitively viewing the low-calorie/high-quantity foods as low in calories or participants ignoring the number. Given the ambiguity of these results, we present Study 5 in the body of the paper and leave this question to future research.

2.6. Facts Study

We first attempted to extend attribute matching to an objective domain by keeping close to valence and presenting participants a series positive statements (e.g., "More than 2 million dogs were adopted from shelters last year") and negative statements (e.g., "There were fewer than 100 murders for every 100,000 citizens in Honduras last year"). The statements were a mix of "more than" and "less than" statements. We manipulated the valence of the decision frame by asking some participants if the statements were true (positive frame) and some if the statements were false (negative frame).

Method

Participants. We decided in advance to recruit 400 participants from Amazon Mechanical Turk, and received 404 (M_{age} = 31.9, 51.0% female).

Materials, procedure, design. Participants viewed five positive statements and five negative statements, in a randomized order. We pretested similar statements with 53 Amazon

Mechanical Turk participants to be perceived as either positive or negative. Half of participants were asked if the statement was true (yes/no); half were asked if the statement was false. (We instructed participants not to cheat and look up the answers, and designed the statements to be difficult to verify online quickly.) After each trial, as in previous studies, participants reported their confidence and consensus estimates. At the conclusion of the study, participants completed a brief attention check and provided their age and gender. (Pre-registration and data: https://osf.io/jrfbq/.)

Results

Fifteen participants (3.7%) failed the attention check and were excluded from analyses.

Confidence. With the remaining participants, we ran an OLS regression, clustering standard errors at the participant level, with *frame*, for the valence of the decision frame (-0.5 = true, 0.5 = false), and *option*, for the valence of the decision's options (0.5 = positive, -0.5 = negative), as well as their interaction. Neither main effect nor the interaction was significant: for the positive statements, participants who were asked if the statements were true were equally confident in their answer (M = 5.53) as the participants asked if the statements were false (M = 5.63), $t(388) = 0.78$, $p = .433$. There was no difference for negative statements either (M_{choose} = 7.22 vs. M_{reject} = 7.48), $t(388) = 0.31$, $p = .759$).

Consensus. We then fit the same model for consensus estimates. Again, neither main effect nor their interaction was significant: for the positive statements, participants who were asked if the statements were true estimated equal consensus in their answer (M = 61.9%) as the participants asked if the statements were false (M = 61.8%), $t(388) = 0.15$, $p = .885$ There was no difference for negative statements either (M_{choose} = 62.9% vs. M_{reject} = 61.9%), $t(388) = 0.92$, p = .357).

2.7. Animal Sizes

An early attempt at attribute matching in objective domains focused on animal size. In this study, we showed participants pairs of animals that were either big (e.g., elephant) or small (e.g., termite). We manipulated this attribute of size in the decision frame by asking some participants which animal in the pair was bigger (big frame) and asking others which animal was smaller (small frame). (Pre-registration and data: https://osf.io/6b8tv/.)

Method

Participants. We decided in advance to recruit 400 participants from Amazon Mechanical Turk, and received 405 (M_{age} = 30.8, 53.6% female).

Materials, procedure, design. Participants viewed five pairs of big animal names and five pairs of small animal names, in a randomized order. Half of participants were asked which animal was bigger; half were asked which animal was smaller. After each pair, participants reported their confidence and consensus estimates as in previous studies. At the conclusion of the study, participants completed a brief attention check, provided their age and gender, and gave the names of any animals we included that they did not recognize.

Results

Twenty-seven participants (6.7%) failed the attention check and were excluded from analyses.

Confidence. With the remaining participants, we ran an OLS regression, clustering standard errors at the participant level, with *frame*, for the valence of the decision frame (-0.5 = which is smaller, 0.5 = which is bigger), and *option*, for the valence of the decision's options (0.5 = big, -0.5 = small), as well as their interaction. We found significant effects of both *frame* (b =

0.12, SE = 0.11, p = .015) and *option* (b = 0.39, SE = 0.06, p < .001). The interaction was significant as well (b = -0.55, SE = 0.12, p < .001), however it was not the crossover interaction we had predicted: for the big animals, there was no difference in reported confidence between the two frames ($M_{bigger?}$ = 6.22, $M_{smaller?}$ = 6.24), $t(376)$ = 0.11, p = .913. For the small animals, participants reported greater confidence when asked which is *bigger* (M = 6.11) than which is smaller (M = 5.57) $t(376)$ = 4.24, p < .001, contrary to our predictions. This in itself was unexpected, but, looking closer at the trial-level effects, the results belied greater inconsistencies: for the big animal pairs, three of the five trended in the predicted direction while the remaining two showed significant effects in the opposite direction; the small animal pairs showed a similar conflicting pattern.

Consensus. We then fit the same model for consensus estimates, and found similar results: both main effects were significant (*frame*: b = 2.73, SE = 0.90, p = .003; *option*: b = 4.18, SE = 0.55, p < .001), but the interaction was, again, not the crossover we predicted (b = -3.20, SE = 1.10, p = .004). There was no difference in consensus estimates for the big animal pairs ($M_{bigger?}$ = 66.5%, $M_{smaller?}$ = 65.4%; $t(376)$ = 1.06, p = .289) and the small animal pairs showed the opposite pattern as we predicted ($M_{bigger?}$ = 63.9%, $M_{smaller}$ = 59.6%; $t(376)$ = 4.11, p < .001). Excluding participants that reported not knowing all the animals included did not meaningfully change the results.

Discussion

The null results from the above two studies initially surprised us. What we believe happened in these two null studies is that the attributes were not nearly salient enough, hence no matching could occur. Additionally, a separate issue may have been that these domains (numeric facts and animal sizes) may have been too obscure for participants to feel comfortable in (note the much lower confidence ratings compared to the first two studies or those in the main text, for example). This concern was one reason why we moved on to (and succeeded with) calorie judgments, a more familiar and accessible domain to most participants.