

Decision

Coherence Shifts in Attribute Evaluations

Douglas G. Lee and Keith J. Holyoak

Online First Publication, June 10, 2021. <http://dx.doi.org/10.1037/dec0000151>

CITATION

Lee, D. G., & Holyoak, K. J. (2021, June 10). Coherence Shifts in Attribute Evaluations. *Decision*. Advance online publication. <http://dx.doi.org/10.1037/dec0000151>

Coherence Shifts in Attribute Evaluations

Douglas G. Lee^{1, 2} and Keith J. Holyoak³

¹ Division of the Humanities and Social Sciences, California Institute of Technology

² Institute of Cognitive Sciences and Technologies, National Research Council of Italy

³ Department of Psychology, University of California, Los Angeles

In five experiments, people repeatedly judged individual options with respect to both overall value and attribute values. When required to choose between two snacks, each differing in two attributes (pleasure and nutrition), people's assessments of value shifted from pre- to post-choice in the direction that spread the alternatives further apart so as to favor the winner, thereby increasing confidence in the choice. This shift was observed not only for ratings of overall value, but also for each of the two individual attributes. The magnitude of the coherence shift increased with choice difficulty as measured by the difference in initial ratings of overall value for the two options, as well as with a measure of attribute disparity (the degree to which individual attributes "disagree" with one another as to which option is superior). In Experiments 2–5, tasks other than explicit choice generated the same qualitative pattern of value changes, confidence, and response time (RT). These findings support the hypothesis that active consideration of options, whether or not explicitly related to value, automatically refines the mental value representations for the options, which in turn allows them to be more precisely distinguished when later included in a value-based choice set.

Keywords: choice, value, coherence shifts, spreading of alternatives, confidence

Supplemental materials: <https://doi.org/10.1037/dec0000151.supp>

The traditional view of multi-attribute decision-making assumes that when people choose between options, they compare them with respect to a set of decision-relevant attributes. According to one theoretical perspective, the attributes of each option are assessed individually, weighted according to their importance, and summed together to map onto an estimate of overall value for each option (e.g., Rieskamp et al., 2006; Tversky & Simonson, 1993). Alternatively, attributes may be assessed and compared separately across options,

with differences at the level of individual attributes being used to choose the preferred option (e.g., Hunt et al., 2014; Tversky, 1969). In either case, attribute values constitute preexisting and stable elements that serve as the basic inputs to the decision process.

An alternative perspective emphasizes the constructive nature of preference (e.g., Ariely et al., 2006; DeKay et al., 2011; Holyoak & Simon, 1999; Payne et al., 1999; Russo et al., 1996; Svenson, 1992; also see Lichtenstein & Slovic,

Douglas G. Lee  <https://orcid.org/0000-0001-5892-8694>

Keith J. Holyoak  <https://orcid.org/0000-0001-8010-6267>

Preparation of this article was supported by NSF Grant BCS-1827374. We thank Dan Simon for comments on an earlier draft of this article.

The data for this study are available in the Open Science Framework at <http://doi.org/10.17605/OSF.IO/X8BPA>

The primary analysis code for this study is available in the Open Science Framework at <http://doi.org/10.17605/OSF.IO/X8BPA>

Douglas G. Lee played a lead role in conceptualization, data curation, formal analysis, investigation, resources, software, visualization, and writing and original draft. Keith J. Holyoak played a lead role in funding acquisition. Douglas G. Lee and Keith J. Holyoak equally contributed in methodology, project administration, validation, and writing, review and editing.

Correspondence concerning this article should be addressed to Douglas G. Lee, Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125, United States. Email: DouglasGLee@gmail.com

2006; Warren et al., 2011). Under this view, subjective assessments of attribute values and attribute importance can shift during the decision process so as to cohere with the emerging decision. These shifts will spread the perceived values of the options, increasing the advantage of the eventual winner relative to its competitor(s) and thereby increasing confidence in the decision. Such spreading of alternatives (SoA), or coherence shifts, have been shown to impact decisions about such everyday matters as choosing a restaurant or an apartment to rent (Russo et al., 1996; Simon et al., 2004, 2008), adjudicating legal disputes (Carlson & Russo, 2001; Holyoak & Simon, 1999; Simon, 2012), and evaluating complex issues with moral implications (Spellman et al., 1993; Holyoak & Powell, 2016; Simon et al., 2015). Cognitive evaluations and emotional reactions can jointly undergo coherence shifts during decision-making (Simon et al., 2015). In contrast to classical cognitive dissonance theory (Festinger, 1957; see Harmon-Jones & Harmon-Jones, 2007, for a review), which claims that value shifts are postdecisional, the construction of preference emphasizes that coherence shifts play a critical role in the predecisional dynamics that drive the selection of an option (Simon & Holyoak, 2002). A number of studies have shown that coherence shifts are observed prior to making a firm decision or public commitment (Holyoak & Simon, 1999; Russo & Chaxel, 2010; Russo et al., 1996, 2008; Simon et al., 2001, 2004).

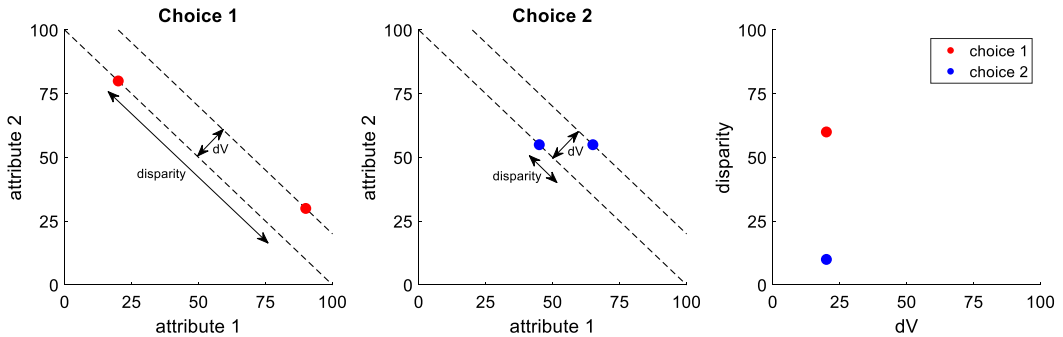
The impact of coherence shifts on decision-making has been firmly established (see Enisman et al., 2021, for a recent meta-analysis). However, less is known about the intra-decisional dynamics that drive attribute reevaluation. The way in which people process information during decision-making may be subject to metacognitive control (e.g., Chaxel, 2015), which may lead to *rational inattention* to different attributes in different choice contexts (Caplin & Dean, 2015; Sims, 2003). The magnitude of shifts in overall evaluations of options tends to increase with the difficulty of the decision (Svenson, 1992), thereby increasing confidence in the choice (Lee & Daunizeau, 2020; Simon et al., 2004; Simon & Spiller, 2016). However, previous studies relating choice difficulty to the magnitude of coherence shifts have not assessed potential changes in the perceived values of those individual attributes that are hypothesized to determine overall value. In the present article, we report a

series of experiments that used a multi-attribute choice paradigm to examine coherence shifts in both overall evaluations and in perceived values of individual attributes. We expected that more difficult choices—those in which two options are initially relatively close in overall value—will trigger larger coherence shifts both for overall value and for individual attributes.

In addition, we examined the impact of a different factor potentially related to choice difficulty. We hypothesized that decision dynamics may also be influenced by differences in attribute composition. Consider the two hypothetical choices situations depicted in Figure 1. Both choices involve two options based on two relevant and equally-weighted attributes, A1 and A2. The first choice is between Option 1 with attribute values of 90 for A1 and 30 for A2, versus Option 2 with values 20 and 80. The second choice is between Option 3 with values 65 and 55, versus Option 4 with values 45 and 55. Under the assumption of equal importance weights for A1 and A2, the two choices are identical at the level of overall value (60 for Option 1 and Option 3, 50 for Option 2 and Option 4). Thus Option 1 is favored for the first choice, and Option 3 for the second choice, by equal amounts. But notice that the disparity between the attributes is much greater for the first choice (Option 1 vs. Option 2) than for the second (Option 3 vs. Option 4).

We propose that the choice with higher disparity will result in larger coherence shifts and thus greater confidence. When choice options are disparate, one option dominates along one attribute dimension, whereas the other dominates along the other dimension. It has been argued that the more dominant is an attribute of a particular option (relative to another option or the average across all other options), the more attention it will receive during the deliberation leading to a decision (Bordalo et al., 2013, 2020). Such increased attention would be expected to enhance the processes that produce spreading of alternatives. An increased coherence shift for high disparity choices would in turn increase choice confidence and decrease response time (RT), as the choice effectively becomes easier (cf. Lee & Daunizeau, 2020).

In addition to addressing the above issues, the present study attempted to identify the core component of the decision-making mechanisms that generate coherence shifts. Coherence shifts evidenced by SoA have typically been assessed in a choice context (hence SoA is commonly

Figure 1*A Schematic Illustration of Orthogonal Components of Choice Difficulty: dV and Disparity*

Note. The left plot illustrates a “high disparity” choice, and the middle plot illustrates a “low disparity” choice. The red and blue dots represent the alternative options for each choice, each plotted according to its measurements on two attribute dimensions. The example assumes equal importance weights for each attribute, so the iso-value curves are represented by parallel lines with slope -1 . The difference in overall value of the options, dV , is the distance between the iso-value curves on which the options lie. Disparity is the distance between the options in the dimension orthogonal to overall value (see Equation 1 for a mathematical formulation). The right plot shows the location of each choice pair in the transformed dV -disparity space. dV = value difference. See the online article for the color version of this figure.

referred to as *choice-induced preference change* (CIPC); Izuma et al., 2010; Voigt et al., 2019). There is evidence, however, that coherence shifts can be produced by meaningful processing of options without introducing an overt choice task (DeKay, 2015; DeKay et al., 2012; Janis & Mann, 1977; Montgomery & Willen, 1999; Russo et al., 1996, 2008; Simon et al., 2001; Svenson, 1992; for reviews see Brownstein, 2003; Russo, 2015). Moreover, it has been suggested that CIPC may actually be an illusion based on a statistical artifact of the rating \rightarrow choice \rightarrow rating (RCR) task design. A study by Chen & Risen (2010) showed that choices between items that are (subjectively) rated as being very close in value will lead to apparent CIPC (on average), even if there is no true cognitive basis for the effect. If the latent mental representations of the option values are assumed to be imprecise (i.e., they have a variance around their expected value), sometimes the initial rating for the to-be chosen item will be drawn from the low tail of its value distribution, and the initial rating for the to-be rejected item will be drawn from the high tail of its value distribution. The final ratings, however, would most likely be closer to the “true” values (i.e., the means of the respective value distributions), due to the laws of probability. In this case, the observed CIPC would have been caused by a statistical artifact of the sampling procedure, not cognitive

reflection. Chen and Risen introduced a new experimental condition to control for this statistical explanation, rating \rightarrow rating \rightarrow choice (RRC). In accord with their statistical hypothesis, they found a significant level of CIPC even in this control condition, supporting their point. However, a number of subsequent studies have shown that CIPC in the standard RCR condition is significantly higher than in the RRC condition, suggesting that choosing between options does indeed cause CIPC beyond what can be explained by statistical considerations (Chammat et al., 2017; Coppin et al., 2014; Lee & Daunizeau, 2020; Salti et al., 2014; Sharot et al., 2012; see Enisman et al., 2021, for a meta-analysis).

It has been suggested that deliberation may induce CIPC as representations of option values are refined (Lee & Daunizeau, in press). However, there has as yet been little systematic investigation of the extent to which different types or different degrees of refinement may take place during tasks other than explicit choice deliberation. We hypothesize that the refinement of value representations is a continuous rather than binary phenomenon, and that different tasks will elicit degrees of coherence shifts that fall along a spectrum. The high end of the spectrum should result when there is explicit deliberation about which option to choose (e.g., the standard RCR); the low end of this spectrum should result when there is no task other than repeated ratings

(e.g., the RRC control). Other tasks may result in intermediate degrees of coherence shifts, in proportion to the amount of value-relevant information processing that takes place. We propose that the core component of decision-making that triggers coherence shifts is active consideration of the options, whether or not such consideration is explicitly related to value. Clearly, the more the task at hand resembles a choice (e.g., assessing similarity of the values of two options; Experiment 2), the more value-related refinement should occur. Tasks that involve comparing options but not explicitly on value (e.g., generic similarity judgments; Experiment 3) should lead to a moderate level of value-related refinement. Tasks that neither focus on value nor require comparison of options (Experiments 4–5) should lead to lower (but still detectable) degrees of refinement. Such a pattern would be consistent with previous work showing that merely repeating isolated evaluations of options causes the evaluations to become more precise (Lee & Coricelli, 2020).

We conducted a series of behavioral experiments in which people judged individual options with respect to both overall value and attribute values, both prior to and after making a choice (Experiment 1) or performing some other task based on the same set of stimuli (Experiments 2–5). In the main experiment (Experiment 1), we assessed how the relationship between the options (in terms of both differences in overall value and in attribute disparity) impacted eventual choice, choice confidence, RT, and shifts in overall value and in attribute values. In the auxiliary experiments (Experiments 2–5), we assessed how the degree of coherence shifts varied as a function of (assumed) value-relevant information processing.

Experiment 1

Experiment 1 examined shifts in overall option values and in attribute values triggered by the choice between two snacks with varying initial values on two attributes, pleasure and nutrition.

Method

Participants

A total of 58 people participated in Experiment 1 (32 female; age: mean = 39 years, $SD = 8$, range 25–50). The data for one participant were corrupt, so we excluded this

participant from analyses.) This sample size was chosen to be comparable to that used in previous studies based on a similar paradigm. All participants were recruited using Amazon Mechanical Turk (MTurk; <https://mturk.com>). All were classified as “masters” by MTurk. They were residents of the United States or Canada, and all were native English speakers. Each participant received a payment of \$7.50 as compensation for approximately 1 hr of time. Our experiments involved de-identified participant data, and protocols were approved by the Institutional Review Board of the University of California, Los Angeles. All participants gave informed consent prior to commencing the experiments.

Materials

The experiments were constructed using the online experimental platform Gorilla (gorilla.sc). The experimental stimuli were drawn from a set of 200 digital images used in a previous study (Lee & Coricelli, 2020), each representing a distinct snack food item. For Experiment 1, we used a subset of 100 stimuli, identical for all participants. We determined this subset by selecting the 100 items for which ratings from the previous study varied the least (within each item) across participants. Figure 2 shows examples of images of snacks, indicating their initial values on the two attributes of pleasure and nutrition (averaged across participants).

Design and Procedure

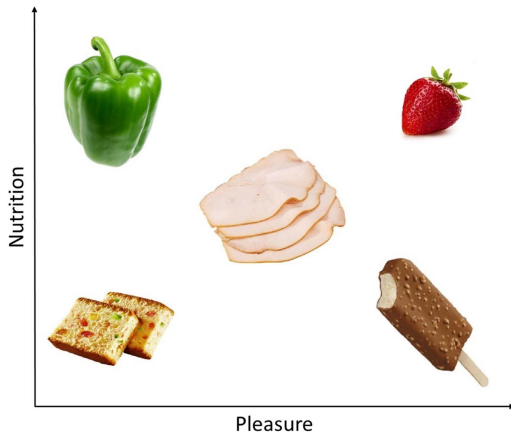
The experiment consisted of a pre-exposure phase, followed by initial value ratings, choice, and final ratings. No time limits were imposed for any of the constituent tasks or for the overall experiment.

In the pre-exposure phase, participants simply observed as all individual items were displayed in a random sequence for 750 ms each. The purpose of the pre-exposure phase was to familiarize participants with the full set of items that they would later evaluate, allowing them to form an impression of the range of subjective values across the item set.

The initial value ratings comprised three sets of ratings: for overall value of the snack, and for the value of each of two individual attributes, nutrition and pleasure. In each rating task, all stimuli were displayed on the screen, one at a time, in a sequence

Figure 2

An Example of Snack Foods Included in the Stimulus Set, Plotted According to Their Respective Ratings Along the Dimensions of Pleasure and Nutrition (Averaged Across Participants)



Note. From lower left to upper right: fruitcake rated low on both pleasure and nutrition; bell pepper rated low on pleasure but high on nutrition; sliced turkey rated medium on both pleasure and nutrition; ice cream bar rated high on pleasure but low on nutrition; strawberry rated high on both pleasure and nutrition. See the online article for the color version of this figure.

randomized across participants. At the onset of each trial, a fixation cross appeared at the center of the screen for 750 ms. Next, an image of a single food item appeared at the center of the screen. For the rating of overall value, participants responded to the question, “How much would you like this as a daily snack?” using a horizontal slider scale. This question was intended to motivate participants to think carefully while assessing the overall subjective quality of each option, as the choice was to determine long-term daily consumption, rather than a “one off” snack. The leftmost end of the scale was labeled “Not at all,” and the rightmost end was labeled “Very much!” The scale appeared to participants to be continuous, and the data was captured in increments of 1 (ranging from 1 to 100). Participants could revise their rating as many times as they liked before finalizing it. Participants clicked the “enter” button to finalize their value rating response and proceed to the next screen. The next trial then began.

The overall value ratings were followed by attribute ratings, which were performed separately for nutrition and pleasure, with the order counter-balanced across participants. One attribute was rated for every alternative (in one task section),

then the other attribute was rated for every alternative (in the next task section). In each attribute rating task, all stimuli were displayed on the screen, one at a time, in a random sequence (randomized across participants and across sections for each participant). The format of this task was identical to that of the overall value rating task, except now participants responded to the question, “How nutritious do you consider this item to be?” or “How pleasurable do you consider this item to be?” For both attribute ratings, the leftmost end of the slider scale was labeled “Very low!” and the rightmost end was labeled “Very high!”

The choice task was then administered. For this task, 50 pairs of stimuli were displayed on the screen, one pair at a time, in a sequence randomized across participants. The pairings of items for each choice trial were created so as to make the choices relatively difficult, as assessed by small differences in value ratings between the two items in a choice pair in a previous study (Lee & Coricelli, 2020). Each individual item occurred in a single choice pair. At the onset of each trial, a fixation cross appeared at the center of the screen for 750 ms. Next, a pair of images of food items appeared on the screen, one left and one right of center. Participants responded to the question, “Which would you prefer as a daily snack?” by clicking on the image of their preferred item. Participants then responded to the question, “How sure are you about your choice?” using a horizontal slider scale. The leftmost end of the scale was labeled “Not at all!” and the rightmost end was labeled “Absolutely!” Participants could revise their confidence report as many times as they liked before finalizing it. Participants clicked the “enter” button to finalize their confidence report and proceed to the next screen.

Finally, participants made final ratings of overall value and attribute values, exactly as in the initial ratings, except with stimuli presented in new random orders. Note that this procedure (randomizing order of individual items) serves to separate the ratings from the prior choice context. Prior to completing these final ratings, participants were instructed not to try to remember their earlier ratings, but rather to simply rate the stimuli as they now evaluated them.

Results and Discussion

Due to the lack of experimental control in online experiments, we anticipated that not all participants

would perform the tasks properly. In particular, we received feedback from some participants that the experiment was rather tedious, due to the repetitive nature required by its design. We therefore performed checks to assess whether participants might have given sloppy responses in the later sections of the experiment. Specifically, we calculated the Spearman correlation between first and last ratings, within participant. We also used generalized linear model (GLM) logistic regression of value ratings on choice to calculate the slope coefficient for each participant (separately using first and last ratings). Using these two measures, we deemed participants with scores outside a cutoff (median $\pm 3 \times$ median average deviation) to be outliers and removed them from our analyses. This procedure resulted in the removal of eight participants for Experiment 1. All analyses were therefore based on the 50 remaining participants.

Unless stated otherwise, all statistical effect sizes reported below were first calculated at the individual level and then tested for significance at the population level. All reported p values represent the probability of non-zero effect sizes, based on standard two-sided t -tests. (See the Supplemental Materials for tables containing detailed statistical summaries for each of the effects reported below.) For every GLM regression analysis reported below, all variables were first converted to z -scores (within participant) before being entered into the models. To assist with both readability and interpretation, we coded all variables such that Option 1 (for each choice) refers to the option with the higher overall value rating in the first phase. We thus define dV (value difference) as the difference in overall value ratings (Option 1 minus Option 2). We define dP (pleasure difference) and dN (nutrition difference) in an analogous manner.

Coherence Shifts in Overall Value

Our primary analysis focus is on coherence shifts, which result in spreading of alternatives (SoA) from initial to final ratings. The choice defines the winning option, and SoA is defined in terms of changes that relatively favor the winner. SoA can be defined for both overall value ratings and for ratings of individual attribute values. For overall value, SoA is defined as the change in overall value for the chosen option from initial to final rating minus the change in

overall value for the unchosen option. In accord with previous work showing coherence shifts over the course of decision-making, we observed a reliable SoA in overall value across all choice trials (cross-participant mean of within-participant median SoA = 2.82, $p < .001$), which is comparable to the magnitude observed in previous studies (Izuma et al., 2010; Lee & Daunizeau, 2020; Voigt et al., 2019. See the Supplemental Materials for more details on SoA.)

We then assessed the relationship between choice difficulty measured by the difference in initial overall value ratings between the options, dV_1 , where the subscript distinguishes the initial ratings (1) from the final ratings (2). As in previous studies (Lee & Coricelli, 2020; Lee & Daunizeau, in press), dV_1 had a reliable negative relationship with SoA and RT, and a reliable positive relationship with choice consistency and choice confidence. We used GLM to regress dV_1 on overall SoA, on RT, on consistency, and on confidence, separately. The beta weights for dV_1 on each dependent variable were significant and in the predicted direction: mean beta for SoA = -0.247 ($p < .001$); mean beta for RT = -0.218 ($p < .001$); mean beta for consistency = 1.634 ($p < .001$); mean beta for confidence = 0.354 ($p < .001$).

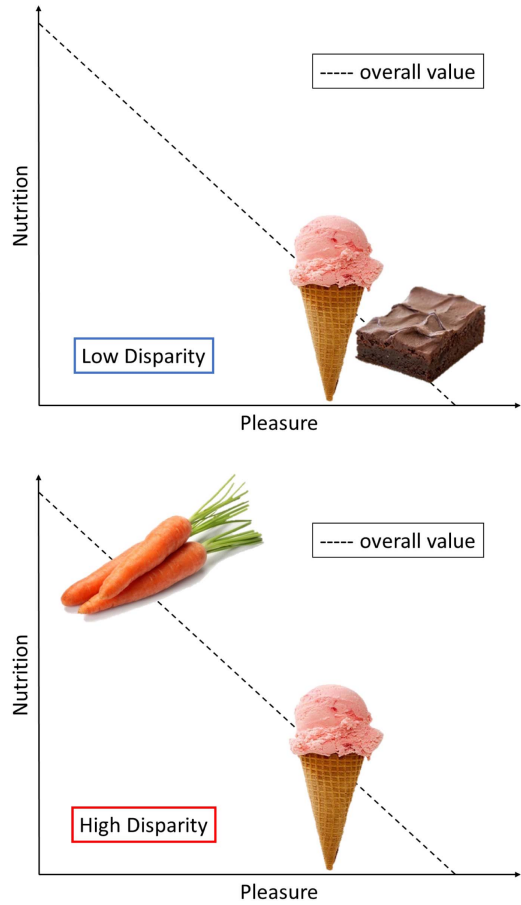
Lee and Daunizeau (2020) proposed that decision-makers refine their value estimates and certainty for the choice options during deliberation, prior to committing to the choice. In support of this hypothesis, these investigators showed that the impact of dV on RT, on consistency, and on confidence is larger when dV is calculated using post-choice ratings (i.e., dV_2) rather than pre-choice ratings. This basic finding (also observed by Simon et al., 2004; Simon & Spiller, 2016) was replicated in our Experiment 1. The mean GLM beta weight for dV_2 as a predictor of RT was -0.2478 ($p < .001$). The magnitude of this beta weight was greater than that obtained using dV_1 ($p = .037$, one-sided t -test). The mean GLM beta weight for dV_2 as a predictor of consistency was 2.7106 ($p < .001$). The magnitude of this beta weight was reliably greater than that obtained using dV_1 ($p < .001$, one-sided t -test). The mean GLM beta weight for dV_2 as a predictor of confidence was 0.3955 ($p < .001$). The magnitude of this beta weight was reliably greater than that obtained using dV_1 ($p = .015$, one-sided t -test).

Lee and Daunizeau (2020) also found that SoA positively influences choice confidence, as would be predicted given that the impact of a coherence shift is to effectively make the choice easier prior to entering a response. In brief, a high value of dV_1 will make the choice an easy one, and confidence will therefore be high. A low dV_1 will make the choice difficult, thereby encouraging deliberation before responding. Deliberation will on average result in a large SoA (which is essentially an increment in dV_1), and therefore will lead to higher confidence. When we included both dV_1 and SoA in a GLM regression model to predict confidence, the cross-participant mean beta weight for dV_1 was 0.397 ($p < .001$) and for SoA was 0.168 ($p < .001$). We also checked for a similar effect of SoA on RT. When we included both dV_1 and SoA in a GLM regression model to predict RT, the cross-participant mean beta weight for dV_1 was -0.241 ($p < .001$) and for SoA was -0.062 ($p = .013$).

Finally, we tested our predictions that attribute disparity would correlate positively with both SoA and choice confidence, and negatively with RT. To explicitly define our disparity variable, we transformed the space representing choice options from attribute space (i.e., a two-dimensional space composed of pleasure and nutrition axes) to a new space in which one dimension is dV_1 and the other is a measure of disparity. Each point in the original attribute space (representing an individual item) necessarily resides on a specific iso-value line (i.e., an imaginary line that connects all points of equal overall value), the slope of which is determined by the (participant-specific) marginal rate of substitution (MRS) of the attributes. We calculated the MRS for each participant as $-b_P/b_N$, where b_P and b_N are the beta weights from the regression of pleasure and nutrition on overall value. The difference in the overall value of the two options being compared is thus the distance between the two iso-value lines on which the options reside (measured in the direction orthogonal to the lines). The disparity measure that we seek is the distance between the two options in the direction parallel to the iso-value lines (see Figure 3 for an illustrative example). This is simply the distance between the scalar projections of the two points (i.e., the location of the options in attribute space) onto the iso-value vector $[-b_P \ b_N]^T$:

Figure 3

Illustration of Two Choice Sets for Snack Foods, One Low Disparity (Top Plot), One High Disparity (Bottom Plot)



Note. As shown by the dashed iso-value lines, all of the available snacks are of comparable overall value (and thus each choice pair is of comparable low dV). However, the two choice pairs are of very different disparity. In the low disparity pair (top), both options score high on pleasure and low on nutrition. In the high disparity pair (bottom), one option scores high on pleasure but low on nutrition, while the other option scores low on pleasure but high on nutrition. dV = value difference. See the online article for the color version of this figure.

$$\text{disparity}_{ij} \triangleq \left| \frac{\begin{bmatrix} P_i N_i \end{bmatrix} * \begin{bmatrix} -b_P \\ b_N \end{bmatrix} - \begin{bmatrix} P_j N_j \end{bmatrix} * \begin{bmatrix} -b_P \\ b_N \end{bmatrix}}{\left\| \begin{bmatrix} -b_P \\ b_N \end{bmatrix} \right\|} \right|,$$

for options i, j

(1)

Using this formal definition of disparity, we performed GLM regressions at the participant level of dV_1 and disparity₁ (where the subscript indicates disparity based on initial ratings; for simplicity we henceforth drop the subscripts for the two options) on SoA, RT, choice consistency, and choice confidence. Across participants, all beta weights were significant and in the predicted direction, though the impact of disparity on RT was not significant (mean dV_1 beta for SoA = -0.344 , $p < .001$; mean disparity₁ beta for SoA = 0.184 , $p < .001$; mean dV_1 beta for RT = -0.195 , $p < .001$; mean disparity₁ beta for RT = -0.037 , $p = .112$ mean dV_1 beta for consistency = 1.558 , $p < .001$; mean disparity₁ beta for consistency = 0.249 , $p = .022$; mean dV_1 beta for confidence = 0.313 , $p < .001$; mean disparity₁ beta for confidence = 0.071 , $p = .026$; see Figure 4).

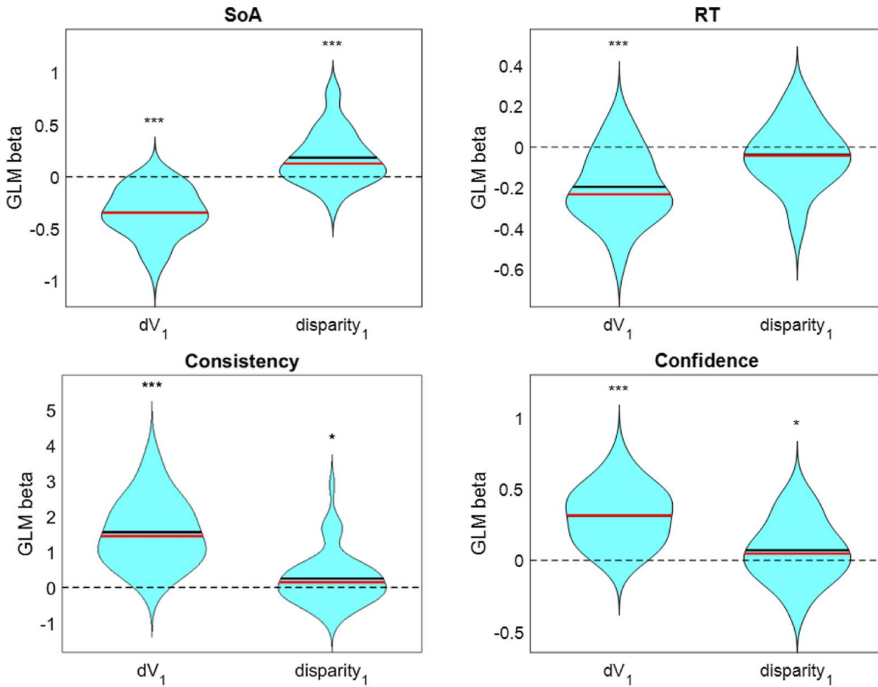
Coherence Shifts in Attribute Values

The next set of analyses focused on coherence shifts at the level of the individual attributes, pleasure and nutrition. SoA values at the attribute level (SoA_P and SoA_N for pleasure and nutrition, respectively) were calculated in an analogous manner as overall SoA—in terms of the magnitude of the shift between the initial and final ratings of attribute values in the direction favoring the winner of the choice. Similarly, we defined the value difference for each attribute separately. For the initial ratings, these differences are termed dP_1 and dN_1 , respectively.

We first checked to see if there was a reliable positive average SoA_P and SoA_N across trials. We observed a reliable SoA_P across all items (cross-participant mean of within-participant median SoA_P = 1.04 , $p = .025$). We observed a smaller

Figure 4

Impact of dV_1 and Disparity₁ on SoA (Top Left), RT (Top Right), Consistency (Bottom Left), and Confidence (Bottom Right)

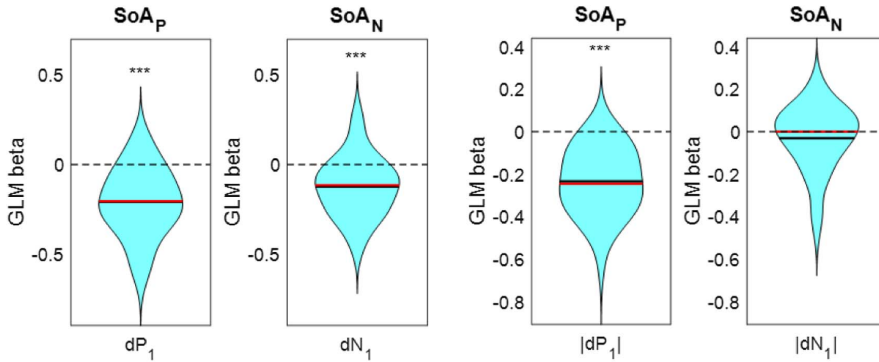


Note. Violin plots represent cross-participant distributions of GLM beta weights; black lines represent cross-participant mean values, red lines represent cross-participant median values. dV = value difference; SoA = spreading of alternatives; RT = response time; GLM = generalized linear model. See the online article for the color version of this figure.

* $p < .05$. ** $p < .01$. *** $p < .001$.

Figure 5

Impact of dP_1 on SoA_P and dN_1 on SoA_N (Left) and of $|dP_1|$ on SoA_P and $|dN_1|$ on SoA_N (Right)



Note. The left panel demonstrates the impact of the relative attribute rating of the chosen option versus the rejected option; the right panel demonstrates the impact of attribute rating similarity, regardless of which item was chosen. dP = pleasure difference; dN = nutrition difference; SoA_N = spreading of alternatives for nutrition; SoA_P = spreading of alternatives for pleasure; GLM = generalized linear model. See the online article for the color version of this figure.

and less reliable SoA_N across all items (cross-participant mean of within-participant median $SoA_N = 0.56$, $p = .108$. See the Supplemental Materials for more details on SoA_P and SoA_N .)

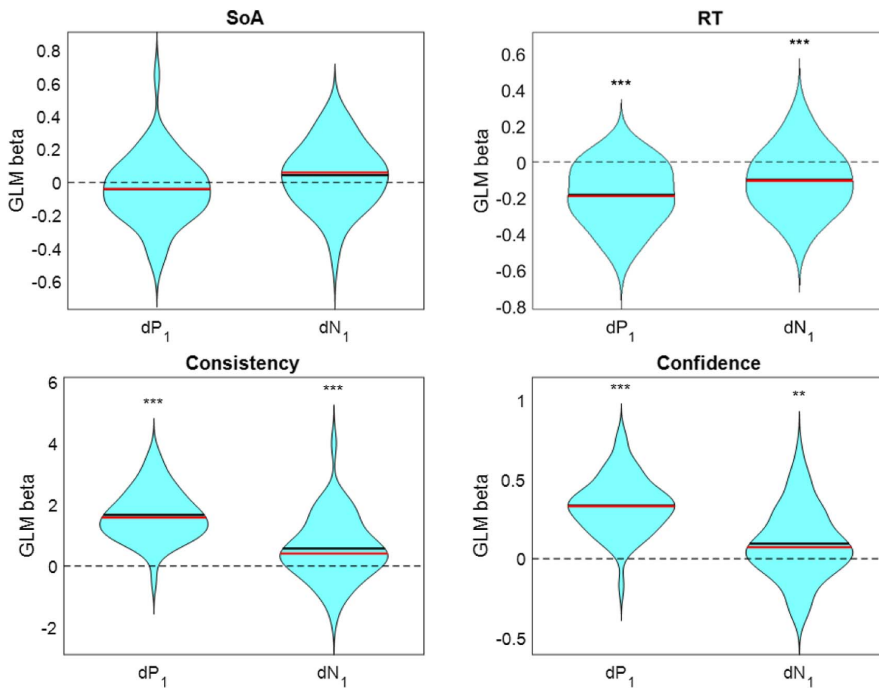
We then tested our predictions that dP_1 would predict SoA_P and that dN_1 would predict SoA_N using GLM regression. Across participants, both beta weights were negative and significant (for dP_1 as a predictor of SoA_P , mean beta = -0.210 , $p < .001$; for dN_1 as a predictor of SoA_N , mean beta = -0.124 , $p < .001$; see Figure 5). The larger beta value for pleasure than for nutrition likely reflects the greater subjective importance of the former attribute. Recall that dP and dN were coded as the attribute ratings for the higher valued option minus the attribute ratings for the lower valued option (based on pre-choice overall value ratings). These results imply that the attribute ratings shifted in favor of the chosen option, and to a larger extent if the chosen option initially rated poorly on that attribute (relative to the rejected option). If this attribute-specific SoA was merely a statistical artifact, we would expect the same regression results when using the absolute value of dP or dN (because the distinction of chosen vs. rejected would no longer be relevant). We repeated the same regressions using $|dP_1|$ and $|dN_1|$ instead of dP_1 and dN_1 . Across participants, the beta weight for $|dP_1|$ was negative and significant (mean beta = -0.234 , $p < .001$), but the beta weight for $|dN_1|$ was not significant (mean beta = -0.031 , $p = .169$; see Figure 5). The fact

that the beta weights for dP_1 and for $|dP_1|$ were similar likely reflects the fact that the pleasure rating for the chosen option was usually higher than for the rejected option (cross-participant median portion of trials = 0.72), whereas this was less frequently true for the nutrition rating (cross-participant median portion of trials = 0.50). We also assessed whether dP_1 and dN_1 would predict overall SoA . Across participants, neither beta weight was significant (dP_1 $p = .156$, dN_1 $p = .105$; see Figure 6). This suggests that overall SoA was not simply an affine combination of SoA_P and SoA_N , and that coherence shifts occurred independently within the specific attribute dimensions.

We then ran GLM regressions of both dP_1 and dN_1 on choice consistency, on choice confidence, and on RT. Across participants, beta weights for consistency were positive and significant for both independent variables (mean dP_1 beta = 1.663 , $p < .001$; mean dN_1 beta = 0.566 , $p < .001$), beta weights for confidence were positive and significant for both independent variables (mean dP_1 beta = 0.338 , $p < .001$; mean dN_1 beta = 0.096 , $p = .004$), and beta weights for RT were negative and significant for both independent variables (mean dP_1 beta = -0.182 , $p < .001$; mean dN_1 beta = -0.100 , $p < .001$; see Figure 6). When post-choice instead of pre-choice attribute ratings were used as predictors of consistency, the beta weight for dP_2 was reliably greater than that for dP_1 ($p = .003$, one-sided t -test), but that for

Figure 6

Impact of Both dP_1 and dN_1 on SoA (Top Left), on RT (Top Right), on Consistency (Bottom Left), and on Confidence (Bottom Right)



Note. Violin plots represent cross-participant distributions of GLM beta weights; black lines represent cross-participant mean values, red lines represent cross-participant median values. dP = pleasure difference; dN = nutrition difference; SoA = spreading of alternatives; RT = response time; GLM = generalized linear model. See the online article for the color version of this figure.

* $p < .05$. ** $p < .01$. *** $p < .001$.

dN_2 did not differ from that for dN_1 ($p = .554$, one-sided t -test). When post-choice attribute ratings were used as predictors of confidence, the beta weight for dP_2 did not differ from that for dP_1 ($p = .754$, one-sided t -test), nor did the beta weight for dN_2 differ from that for dN_1 ($p = .934$, one-sided t -test). When post-choice ratings were used as predictors of RT, the beta magnitude for dP_2 was reliably greater than that for dP_1 ($p = .072$, one-sided t -test), but that for dN_2 did not differ from that for dN_1 ($p = .961$, one-sided t -test).

Experiments 2–5

Experiments 2–5 were designed to test the hypothesis that tasks other than explicit choice are also able to trigger coherence shifts. The general materials and design were basically the same as those of Experiment 1, except that ratings were obtained in three phases: initial ratings,

ratings after a non-choice task, and final ratings after a choice. As in Experiment 1, the final choice defined the winning option with respect to which all SoA measures were defined.

Method

Participants

A total of 63 people participated in Experiment 2 (31 female; age: mean = 41 years, $SD = 8$, range 26–50). A total of 68 people participated in Experiment 3 (45 female; age: mean = 42 years, $SD = 8$, range 19–50). A total of 67 people participated in Experiment 4 (37 female; age: mean = 39 years, $SD = 8$, range 25–50). A total of 69 people participated in Experiment 5 (41 female; age: mean = 41 years, $SD = 7$, range 27–50). All participants were MTurk “masters,” residents of the United States or Canada, and native English speakers. Each participant received a payment of

either \$7.50 or \$9 (increased for Experiments 4–5) as compensation for approximately 1 hr of time.

As in Experiment 1, we filtered our data for participants who might have performed the tasks carelessly. For each experiment, we calculated the Spearman correlation between first and last ratings, within participant, and also used GLM logistic regression of value ratings on choice to calculate the slope coefficient for each participant (separately using first and last ratings). Using these two measures, we deemed anyone with scores outside a cutoff (median + 3 × median average deviation) to be an outlier and removed them from our analyses. After removing outliers, the number of participants whose data was used in analyses was 44 for Experiment 2, 48 for Experiment 3, 51 for Experiment 4, and 55 for Experiment 5.

Materials

For Experiments 2–5, we selected 60 stimuli (identical for all participants) from the full set of snacks (each presented as a digital image). The experimental set consisted of the 30 choice pairs (identical for all participants) that generated the highest overall SoA across participants in Experiment 1.

Design and Procedure

The basic design was very similar to that of Experiment 1, but with three phases of ratings (rather than two). After pre-exposure to all individual snacks, participants gave their initial ratings (Phase 1) of overall value and of attribute values. This was followed by a non-choice task related to the 60 items, after which another set of ratings was obtained (Phase 2). The choice task was then administered in the same way as in Experiment 1, and a final set of ratings was obtained (Phase 3). Prior to administering the choice task, participants were not informed that any task would require a choice between options.

The pre-exposure, rating, and choice tasks were identical in format to those in Experiment 1. Across Experiments 2–5, the non-choice task varied from more to less “choice-like.” In Experiments 2–4, the 30 pairs of stimuli were displayed on the screen, one pair at a time, in a sequence randomized across participants (hence the two items were visible for potential comparison). These item pairs were the same as those that would be used in the choice task, except for the presentation sequence. Each individual item occurred in a single comparison pair.

At the onset of each trial, a fixation cross appeared at the center of the screen for 750 ms. Next, a pair of images of food items appeared on the screen, one left and one right of center. In Experiment 2, participants responded to the question, “How similarly would you like these as daily snacks?” using a dropdown menu. The response choices were, “I like them {totally equal, very similar, fairly similar, slightly similar, totally different} amounts.” In Experiment 3, participants responded to the question, “How similar are these?” using a dropdown menu. The response choices were {totally similar, very similar, fairly similar, slightly similar, totally different}. Thus, whereas the similarity judgment made in Experiment 2 referred explicitly to “liking” of snacks, the similarity judgment in Experiment 3 made no reference to any sort of value. In Experiment 4, a green arrow appeared above one of the images on each trial (randomized between left and right), indicating which item the participants were to assess. Participants responded to the question, “When would you prefer to eat this?” using a dropdown menu. The response choices were {morning, afternoon, evening}. This task did not require participants to compare the two items in a pair, but did not prevent them from making a comparison.

In Experiment 5, in contrast to Experiments 1–4, the options were presented one at a time (rather than as pairs), so that a direct comparison of two items was impossible. Participants responded to the identical question as in Experiment 4 (i.e., time-of-day preference for a single snack). In all experiments, participants could revise their response as many times as they liked before finalizing it by clicking the “enter” button and proceeding to the next screen. No time limits were imposed for any of the constituent tasks or for the overall experiment.

Results and Discussion

Experiments 2–5 included three phases of ratings. We will refer to measures of value difference (overall and for individual attributes) calculated using the initial round of ratings (prior to the non-choice task) with the subscript 1 (e.g., dV_1), those based on the intermediate round of ratings (following the non-choice task but prior to the choice task) with the subscript 2, and those calculated using the final round of ratings (following the choice task) with the subscript 3. For SoA, subscript 1 is used for those measures calculated from

pre- to post-non-choice (i.e., from Phase 1 to Phase 2). The subscript 2 is used for those measures calculated from post-non-choice to post-choice (i.e., from Phase 2 to Phase 3). For all measures of SoA, the winning choice is defined by the option eventually selected during the choice task. Critically, SoA_1 values reflect coherence shifts that occurred *before* the choice task was even administered. Positive values of SoA_1 , $SoAP_1$, and $SoAN_1$ will thus indicate changes in the direction of the eventual winner triggered by a task that did not itself require making a choice. Similarly, any impact of non-choice tasks on eventual choice confidence precedes the posing of a choice task.

Coherence Shifts in Overall Value During Non-Choice Tasks

Our main goal in Experiments 2–5 was to assess whether each of the effects related to choice that were observed in Experiment 1 would also be triggered by our novel non-choice task design. Indeed, we found a reliable positive level of SoA_1 when we examined the change in overall value ratings before and after the non-choice tasks in each of Experiments 2–5. As anticipated, the magnitude of SoA_1 in these experiments was largest in Experiment 2 (mean of median $SoA_1 = 3.19$, $p < .001$), followed by Experiment 3 (mean of median $SoA_1 = 2.22$, $p = .019$), followed by Experiment 4 (mean of median $SoA_1 = 1.73$, $p = .008$), with Experiment 5 being the lowest (mean of median $SoA_1 = 1.18$, $p = .068$). These results are summarized in Table 1 below, which also includes the comparable SoA_1 value for Experiment 1 (based only on the same 30 choice pairs that were also included in Experiments 2–5). Note the clear decreasing pattern of SoA_1 as the non-choice task moves from more to less “choice-like.” We confirmed that the apparent gradient of

SoA_1 as a function of task was statistically significant using an ANOVA with task (i.e., experiment) as a factor. The overall effect of task was significant, $F(4, 243) = 3.08$, $p = .017$, as was a planned comparison examining the apparent linear trend in the magnitude of the SoA_1 effect, $t(246) = 3.48$, $p < .001$.

The cross-experiment gradient of SoA_1 supports our claim that a major cause of SoA is cognitive contemplation of the options, rather than it being either a mere statistical artifact or triggered only by an explicit choice (as cognitive dissonance theorists would hold). To provide additional support for our claim, we ran GLM regressions of dV_1 , dV_2 , and dV_3 (separately) on eventual choice consistency, on choice confidence, and on RT. All beta weights were significant and in the predicted direction (see Table 2; $p < .001$ for all beta weights). More interestingly, in each experiment the magnitude of the beta weights for dV on each of the dependent variables increased from Phase 1 to 2 and from Phase 2 to 3. Not all of the differences in beta weights were statistically significant (see Figure 7), but the trend was systematic and robust across dependent variables and across experiments (see Table 2 and Figure 7). Thus, value differences after coherence shifts (both between Phases 1 and 2, and between Phases 2 and 3) were better predictors of choice consistency, choice confidence, and RT than were the initial value differences. Notably, this pattern would not be expected if the coherence shifts were due to post-choice cognitive dissonance resolution or were solely a statistical artifact of repeated ratings. Using ANOVA analyses with rating session as dependent measure, we confirmed that the increase in magnitude of the regression weights as a function of rating session was statistically significant for choice confidence, $F(2, 245) = 16.02$, $p < .001$, and RT, $F(2, 245) = 3.06$, $p = .048$, but not for choice consistency, $F < 1$.

The results presented above suggest that ratings provided later in the experiment are more in line with the “true” evaluations revealed by the explicit choice task. This increase in precision should be closely linked to SoA, since SoA_1 and SoA_2 are basically $dV_2 - dV_1$ and $dV_3 - dV_2$, respectively (although they are not mathematically identical, because preference reversals can increase the magnitude of SoA). To confirm the relationship between SoA and the dependent variables of interest, we performed GLM regressions of dV_1 ,

Table 1

Cross-Participant Mean of Within-Participant Median SoA_1 Across Tasks in Five Experiments

Experiment #	SoA_1
1 (choice)	4.200
2 (similarity-value)	3.193
3 (similarity-general)	2.219
4 (time of day-pairs)	1.726
5 (time of day-singles)	1.182

Note. SoA = spreading of alternatives.

Table 2*Impact of dV on Consistency, Confidence, and Response Time (Experiments 1–5)*

Choice consistency	β_{dv1_Ch}	β_{dv2_Ch}	β_{dv3_Ch}
Experiment 1 (choice)	2.861	4.725	—
Experiment 2 (similarity-value)	3.810	6.107	5.924
Experiment 3 (similarity-general)	2.462	2.971	4.385
Experiment 4 (time of day-pairs)	2.698	4.231	4.549
Experiment 5 (time of day-singles)	3.324	4.010	4.492
Choice confidence	β_{dv1_CC}	β_{dv2_CC}	β_{dv3_CC}
Experiment 1 (choice)	0.337	0.418	—
Experiment 2 (similarity-value)	0.393	0.484	0.485
Experiment 3 (similarity-general)	0.391	0.443	0.488
Experiment 4 (time of day-pairs)	0.393	0.451	0.478
Experiment 5 (time of day-singles)	0.355	0.413	0.479
Response time	β_{dv1_RT}	β_{dv2_RT}	β_{dv3_RT}
Experiment 1 (choice)	−0.235	−0.266	—
Experiment 2 (similarity-value)	−0.263	−0.299	−0.303
Experiment 3 (similarity-general)	−0.264	−0.287	−0.316
Experiment 4 (time of day-pairs)	−0.271	−0.301	−0.297
Experiment 5 (time of day-singles)	−0.222	−0.244	−0.282

Note. Regression weights for Experiment 1 are based only on the same 30 choice pairs that were also included in Experiments 2–5. For choice consistency, two extreme outliers in Experiment 1 were excluded.

SoA₁, and SoA₂ on choice confidence and on RT.¹ Across all experiments, all beta weights were significant and in the expected direction (positive for confidence, negative for RT; see Figure 8).

Finally, we verified that our additional regression analyses from Experiment 1 replicated across each of the other experiments. Specifically, we ran GLM regressions of dV_1 and SoA₁ on RT and on confidence. We also ran GLM regressions of dV_1 and disparity₁ on SoA₁, on RT, and on confidence. Across participants, all beta weights were significant and in the predicted direction (see Supplemental Materials, section Statistical Summary).

In addition to the analyses described above, we examined the change in individual attribute ratings before and after the non-choice tasks in each of Experiments 2–5. We found a reliable positive effect (cross-participant mean of within-participant median) for SoAP₁ in each experiment, but not for SoAN₁. Furthermore, we did not find evidence of a gradient of either SoAP₁ or SoAN₁ across the experiments, as we did with SoA₁. Thus, the changes in SoA for individual attributes proved to be less sensitive measures than overall SoA.

We also examined the change in overall value ratings before and after the eventual choice task in each of Experiments 2–5 (i.e., SoA₂). A positive effect was observed in each experiment, but it was only statistically significant in two of the four

experiments, and there was no discernable pattern across experiments (Experiment 2: mean of median SoA₂ = 0.614, $p = .339$; Experiment 3: mean of median SoA₂ = 2.583, $p < .001$; Experiment 4: mean of median SoA₂ = 0.637, $p = .190$; Experiment 5: mean of median SoA₂ = 1.809, $p = .003$). Previous studies have also found that the major shift in SoA occurs from the first to the second ratings, rather than thereafter (e.g., Holyoak & Simon, 1999).

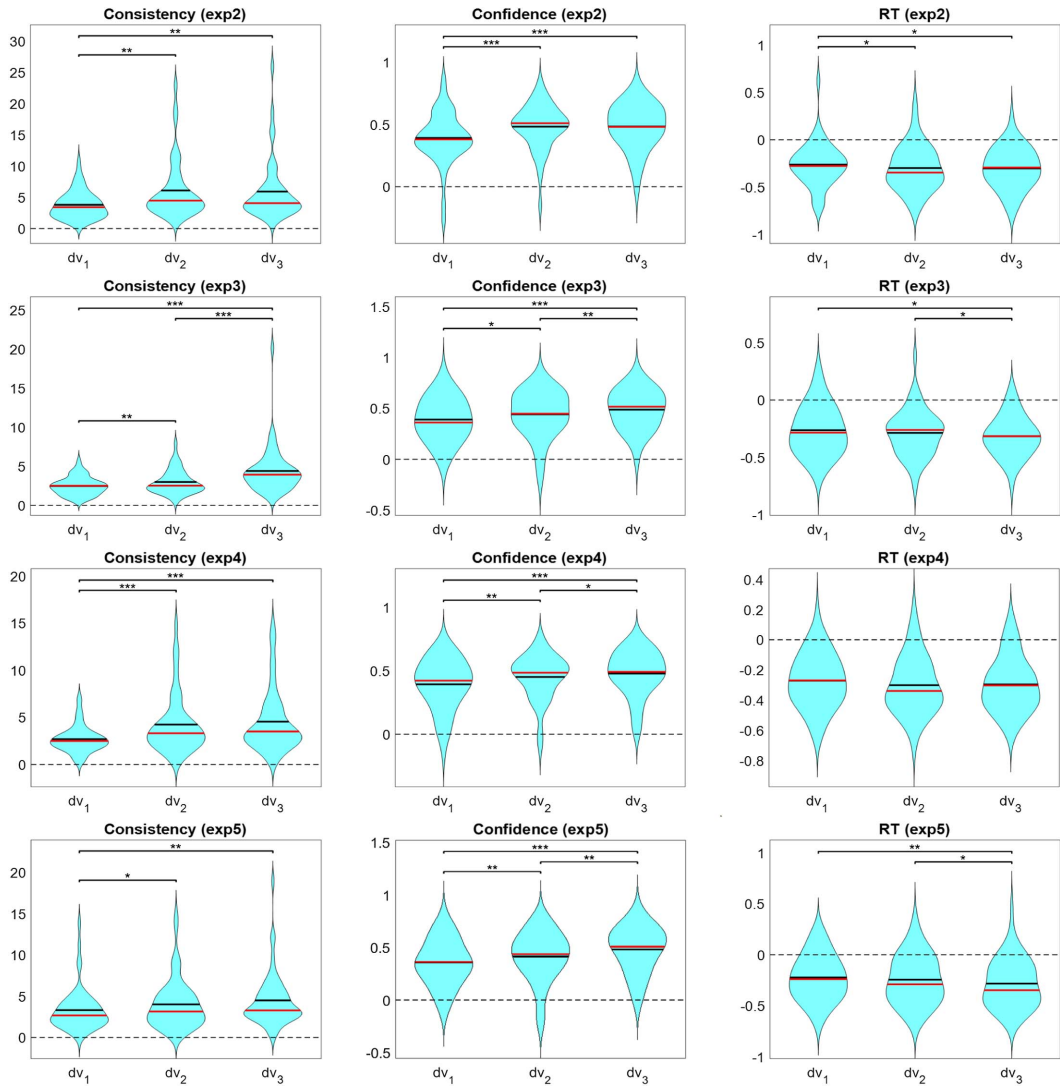
General Discussion

The present study provides a number of important findings concerning the role of coherence shifts in the construction of preferences. When required to choose between two options (snacks), each differing in two attributes that determine value (pleasure and nutrition), people's assessments of value shifted from pre- to post-choice in the direction that spread the alternatives further

¹ We do not include consistency here, because some choices would be classified as preference reversals (i.e., dV_1 and dV_3 would have opposite signs). These cannot necessarily be classified as errors, however, as value refinements might have caused decision-makers to change their mind about their preferences on these trials. For such trials, choice consistency would be uncorrelated or anti-correlated with confidence and RT, even though all three variables might be influenced by SoA.

Figure 7

Impact of dV on Consistency (Left Panels), on Confidence (Middle Panels), and on RT (Right Panels), Separately for Each Round of Ratings



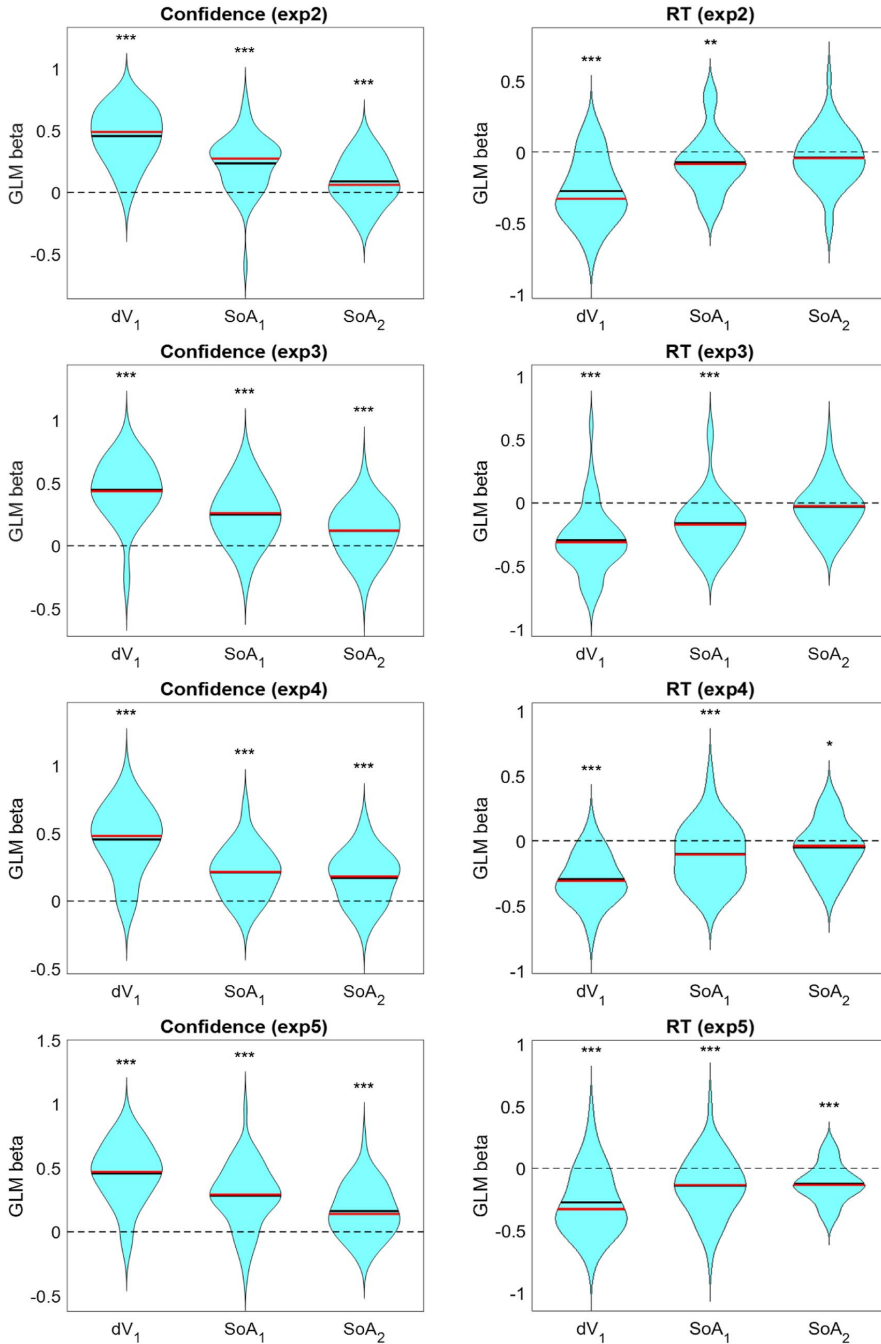
Note. Each row of plots is for a different experiment. Violin plots represent cross-participant distributions of GLM beta weights; black lines represent cross-participant mean values, red lines represent cross-participant median values. dV = value difference; RT = response time; exp = experiment; GLM = generalized linear model. See the online article for the color version of this figure.

* $p < .05$. ** $p < .01$. *** $p < .001$.

apart so as to favor the winner, thereby increasing confidence in the choice. This shift was observed not only for ratings of overall value, but also for each of the two individual attributes that determined value. Moreover, the magnitude of the coherence shift increased with the difficulty of the choice as measured by the difference in initial ratings of overall value for the two options.

Coherence shifts were in turn predictive of increased choice confidence and decreased RT, with each of these dependent variables being more accurately predicted by value difference measured by ratings obtained after the choice.

We also found that coherence shifts are predicted not only by the overall value difference between options, but also by the pattern of

Figure 8*Impact of dV , SoA_1 , and SoA_2 on Confidence (Left Panels) and on RT (Right Panels)*

Note. Each row of plots is for a different experiment. Violin plots represent cross-participant distributions of GLM beta weights; black lines represent cross-participant mean values, red lines represent cross-participant median values. dV = value difference; SoA = spreading of alternatives; RT = response time; exp = experiment; GLM = generalized linear model. See the online article for the color version of this figure.

* $p < .05$. ** $p < .01$. *** $p < .001$.

attribute composition across options (which we believe is a novel contribution to the literature). We introduced a formal measure of attribute *disparity* (Equation 1), reflecting the degree to which individual attributes “disagree” with one another as to which option is superior. After accounting for the impact of overall value difference, the magnitude of the coherence shift was also positively correlated with disparity. Moreover, coherence shifts associated with disparity increased confidence and decreased RT in the eventual choice.

Experiments 2–5 provided evidence that coherence shifts are driven by refinements of the mental representations of the options, and not merely by post-choice adjustments (e.g., cognitive dissonance reduction) or statistical noise (i.e., a regression to the mean effect across repeated evaluations). Specifically, Experiment 2 demonstrated that an active comparison of option values is likely the core component of the decision process that drives the coherence shifts typically observed in a choice task. To isolate the impact of the comparison process, Experiment 2 used a three-phase design in which value ratings were obtained at the outset, after a value comparison task, and finally after an actual choice task. Critically, the comparison task was administered prior to informing the participants that any choice would be required. Instead, they were simply asked to rate the similarity of values for two snacks. This comparison task generated the same qualitative pattern of changes in value ratings and in consistency, confidence, and RT in the eventual choice, even though the similarity judgments entirely preceded the actual choice task. These findings support the hypothesis that active comparison of values is a core component of the decision process during which perceived values are shifted so as to more sharply distinguish the items being compared.

Moreover, Experiments 3–5 demonstrated that an active comparison of option values is not even necessary to cause coherence shifts. Specifically, Experiment 3 replaced similarity judgments based on value comparison (Experiment 2) with a more generic similarity judgment not explicitly linked to value. While it is possible that some participants compared the snacks in terms of value, it is likely that non-value aspects (e.g., size, shape, color) were also considered. The generic similarity task nevertheless triggered

coherence shifts in the eventual choices (between the same pairs of items for which similarity judgments had been made). This finding suggests that any additional processing directed at the options on offer, even if not explicitly cued toward value, leads to a refinement of the mental representations of the options. This refinement automatically brings the benefit of more accurate valuations when solicited in subsequent rating tasks.

In Experiment 4, we replaced the comparison task with an even more generic task. Here, we presented the same pairs of snacks as in the eventual choice task, but asked participants to ignore one option while deciding at what time of day they would prefer to eat the cued option. Notably, this task requires absolutely no comparison between the options. Nevertheless, it appears to have caused some degree of coherence shifts (though less than in the explicit comparison tasks). This finding suggests that when presented with pairs of categorically similar options (in this case, snacks), the brain may automatically perform some initial rough comparison of their values, even when value is not task-relevant and the overt task concerns just one of the paired items. This interpretation is consistent with previous work showing that the brain automatically encodes the relative value of options even when the task is unrelated to value (Grueschow et al., 2015).

In Experiment 5, we used the same time-of-day judgment as in Experiment 4 but with displays showing a single item at a time, thus precluding even the possibility of a comparison between options while performing the task. We still observed some positive level of coherence shifts in Experiment 5 (though with a lower magnitude than in any of the other experiments). This finding suggests that consideration of individual snacks in isolation (rather than in contrast to other snacks) can also lead to refinement of their mental representation. This result is consistent with previous work showing that the brain automatically encodes value of individual options even when the task is unrelated to value (Lebreton et al., 2009). The automatic value signal for a particular option may adjust the latent value estimate in the mind of the decision-maker, altering the probability that that option will later be chosen in a subsequent choice task.

Taken together, our results cast serious doubt on several alternative explanations of observed coherence shifts. Accounts based on post-choice

resolution of cognitive dissonance are unable to explain why we observed coherence shifts in a variety of non-choice situations. Furthermore, dissonance reduction cannot explain the observed gradient in coherence shifts across the different types of tasks. Our data also rule out suggestions that coherence shifts are solely due to statistical artifacts. Under such an account, the value representations of options will not change as a function of intervening tasks. Accordingly, any observed relationship between value difference (dV) and choice behavior (consistency, confidence, RT) should be equally well predicted using ratings from any of the experimental phases. The fact that we observed a clear increase in the explanatory power of dV on all dependent variables from the first to second to third phase makes it highly unlikely that the evolution of dV across phases (via SoA_1 and SoA_2) was mere statistical noise.

Lee and Daunizeau (in press) have recently proposed a computational account of why and how coherence shifts can be generated during intra-decisional processing. According to their metacognitive control of decision-making (MCD) model, decision-makers should exert mental effort only to the degree necessary to reach some subjective threshold of confidence that the chosen option is indeed the best of the candidate set. People process additional information about choice options up until the point at which the options are sufficiently distinguishable to provide a satisfactory level of confidence in the emerging choice. The MCD model can account for empirical evidence that people change their subjective value estimates when required to choose between options that they initially estimated to have similar values, and that this change in value correlates with both choice confidence and RT. Notably, while the authors presented this model specifically in the context of active deliberation during an explicit choice, it could likely be extended to account for the impact of information processing in other tasks (as in Experiments 2–5). For example, contemplation during any task may lead decision-makers closer to some sort of “ground truth” (e.g., if they consider more attributes that they previously ignored), leading to ratings that are both more accurate and more precise (and thus the option with the “true” higher value would be more likely to be chosen in a subsequent choice task). Beyond that, a head-to-

head comparison of options (e.g., during a choice task) may magnify the spreading effect, because the relative values of the attributes (between options) may be more salient than the values of each option evaluated in isolation.

The behavioral evidence concerning intra-decisional processing provided by the present study fits well with the picture arising from studies of the emergence of decisions at the neural level. It has been shown that the brain computes attribute values separately and then integrates them before making a choice (Lim et al., 2013). Other neural evidence indicates that the assessment of individual attributes, as well as their relative weights on the final choice, evolve during the decision process (Hunt et al., 2014). Löffler et al. (2020; also Voigt et al., 2019) have provided compelling neuroimaging evidence that shifts in valuation that support the eventual choice occur prior to the choice response. Future studies using multiple methodologies will hopefully provide a deeper understanding of the timing of coherence shifts at the attribute level. In particular, it would be informative to examine how activity patterns in key brain regions vary across different sorts of choice or non-choice tasks, in which the refinement of value representations seems to take place to varying degrees.

It has been proposed that CIPC is mediated by the strength of memory encoding for individual options. Several studies (Bakkour, et al., 2017; Botvinik-Nezer et al., 2019; Schonberg et al., 2014) showed that enhanced processing of snack food options (based on response cueing for specific options during sequential passive viewing) led to better option recall for those options, and that (regardless of cueing) better recall was associated with a higher probability of being chosen (when paired with alternatives of similar subjective value). This suggests that the more precise value representations associated with cued (or otherwise better remembered) options enhanced the apparent value of those options. Previous work has also shown that CIPC only occurs for options that are remembered well (Chammat et al., 2017). Other previous work has demonstrated stronger memory encoding for options of higher value (Miendlarzewska et al., 2016; Shohamy & Adcock, 2010), which could play a causal role in the spreading of alternatives phenomenon (higher valued options would be more likely to be chosen, and also more precisely encoded).

Schonberg and Katz (2020) proposed a brain network potentially responsible for CIPC, based on selective attention. Further studies are needed to better understand the relationship between value, memory, and preference change.

Another useful direction for future research would be to collect self-reported estimates of attribute-specific importance weights, rather than relying on statistical regression to infer these weights. Recall that in the present experiments, final ratings were obtained using all individual items after completing all choices between paired options. Each individual item occurred in only one choice trial; hence any value shift for an item can be attributed to that one specific choice. However, the two basic attributes were likely assessed on all choice trials, and hence their relative importance would have been pushed in different directions on different trials (presumably canceling out any coherence shifts due to trial-by-trial changes in attribute importance weights). A different paradigm would be required to meaningfully assess coherence shifts in attribute importance. Such a paradigm could also be used to assess whether the magnitude of coherence shifts along a particular attribute dimension is systematically correlated with the self-reported importance of that attribute.

References

- Ariely, D., Loewenstein, G., & Prelec, D. (2006). Tom Sawyer and the construction of value. *Journal of Economic Behavior & Organization*, 60, 1–10.
- Bakkour, A., Lewis-Peacock, J. A., Poldrack, R. A., & Schonberg, T. (2017). Neural mechanisms of cue-approach training. *NeuroImage*, 151, 92–104.
- Bordalo, P., Gennaioli, N., & Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy*, 121(5), 803–843.
- Bordalo, P., Gennaioli, N., & Shleifer, A. (2020). Memory, attention, and choice. *The Quarterly Journal of Economics*, 135(3), 1399–1442.
- Botvinik-Nezer, R., Bakkour, A., Salomon, T., Shohamy, D., & Schonberg, T. (2019). Memory for individual items is related to non-reinforced preference change. *bioRxiv*, Article 621292. <https://doi.org/10.1101/621292>
- Brownstein, A. (2003). Biased pre-decision processing. *Psychological Bulletin*, 129, 545–568.
- Caplin, A., & Dean, M. (2015). Revealed preference, rational inattention, and costly information acquisition. *The American Economic Review*, 105(7), 2183–2203.
- Carlson, K. A., & Russo, J. E. (2001). Biased interpretation of evidence by mock jurors. *Journal of Experimental Psychology: Applied*, 7, 91–103.
- Chammatt, M., El Karoui, I., Allali, S., Hagège, J., Lehongre, K., Hasboun, D., Baulac, M., Epelbaum, S., Michon, A., Dubois, B., Navarro, V., Salti, M., & Naccache, L. (2017). Cognitive dissonance resolution depends on episodic memory. *Scientific Reports*, 7(1), 1–10.
- Chaxel, A.-S. (2015). The impact of a relational mindset on information distortion. *Journal of Experimental Social Psychology*, 60, 1–7.
- Chen, M. K., & Risen, J. L. (2010). How choice affects and reflects preferences: Revisiting the free-choice paradigm. *Journal of Personality and Social Psychology*, 99(4), 573–594.
- Coppin, G., Delplanque, S., Bernard, C., Cekic, S., Porcherot, C., Cayeux, I., & Sander, D. (2014). Choice both affects and reflects preferences. *Quarterly Journal of Experimental Psychology*, 67(7), 1415–1427.
- DeKay, M. L. (2015). Predecisional information distortion and the self-fulfilling prophecy of early preferences in choice. *Current Directions in Psychological Science*, 24(5), 405–411.
- DeKay, M. L., Stone, E. R., & Miller, S. A. (2011). Leader-driven distortion of probability and payoff information affects choices between risky prospects. *Journal of Behavioral Decision Making*, 24(4), 394–411.
- DeKay, M. L., Stone, E. R., & Sorenson, C. M. (2012). Sizing up information distortion: Quantifying its effect on the subjective values of choice options. *Psychonomic Bulletin & Review*, 19(2), 349–356.
- Enisman, M., Shpitzer, H., & Kleiman, T. (2021). Choice changes preferences, not merely reflects them: A meta-analysis of the artifact-free free-choice paradigm. *Journal of Personality and Social Psychology*, 120(1), 16–29.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Row, Peterson.
- Grueschow, M., Polania, R., Hare, T. A., & Ruff, C. C. (2015). Automatic versus choice-dependent value representations in the human brain. *Neuron*, 85(4), 874–885.
- Harmon-Jones, E., & Harmon-Jones, C. (2007). Cognitive dissonance theory after 50 years of development. *Zeitschrift für Sozialpsychologie*, 38(1), 7–16.
- Holyoak, K. J., & Powell, D. (2016). Deontological coherence: A framework for commonsense moral reasoning. *Psychological Bulletin*, 142, 1179–1203.
- Holyoak, K. J., & Simon, D. (1999). Bidirectional reasoning in decision making by constraint satisfaction. *Journal of Experimental Psychology: General*, 128, 3–31.
- Hunt, L. T., Dolan, R. J., & Behrens, T. E. (2014). Hierarchical competitions subserving multi-attribute choice. *Nature Neuroscience*, 17(11), 1613–1622.

- Izuma, K., Matsumoto, M., Murayama, K., Samejima, K., Sadato, N., & Matsumoto, K. (2010). Neural correlates of cognitive dissonance and choice-induced preference change. *Proceedings of the National Academy of Sciences of the United States of America*, 107(51), 22014–22019.
- Janis, I. L., & Mann, L. (1977). *Decision making: A psychological analysis of conflict, choice, and commitment*. Free Press.
- Lebreton, M., Jorge, S., Michel, V., Thirion, B., & Pessiglione, M. (2009). An automatic valuation system in the human brain: Evidence from functional neuroimaging. *Neuron*, 64(3), 431–439.
- Lee, D., & Coricelli, G. (2020). An empirical test of the role of value certainty in decision making. *Frontiers in Psychology*, 11, Article 574473.
- Lee, D., & Daunizeau, J. (2020). Choosing what we like vs. liking what we choose: How choice-induced preference change might actually be instrumental to decision-making. *PLOS ONE*, 15(5), Article e0231081.
- Lee, D. G., & Daunizeau, J. (2021). Trading mental effort for confidence in the metacognitive control of value-based decision-making. *eLife*, 10, Article e63282.
- Lichtenstein, S., & Slovic, P. (Eds.). (2006). *The construction of preference*. Cambridge University Press.
- Lim, S. L., O'Doherty, J. P., & Rangel, A. (2013). Stimulus value signals in ventromedial PFC reflect the integration of attribute value signals computed in fusiform gyrus and posterior superior temporal gyrus. *The Journal of Neuroscience*, 33(20), 8729–8741.
- Löffler, A., Haggard, P., & Bode, S. (2020). Decoding changes of mind in voluntary action—Dynamics of intentional choice representations. *Cerebral Cortex*, 30(3), 1199–1212.
- Miendlarzewska, E. A., Bavelier, D., & Schwartz, S. (2016). Influence of reward motivation on human declarative memory. *Neuroscience and Biobehavioral Reviews*, 61, 156–176.
- Montgomery, H., & Willen, H. (1999). Decision making and action: The search for a good structure. In P. Jusslin & H. Montgomery (Eds.), *Judgment and decision making: Neo-Brunswikian and process-tracing approaches* (pp. 147–173). Erlbaum.
- Payne, J. W., Bettman, J. R., & Schkade, D. A. (1999). Measuring constructed preferences: Towards a building code. *Journal of Risk and Uncertainty*, 19, 243–270.
- Rieskamp, J., Busemeyer, J. R., & Mellers, B. A. (2006). Extending the bounds of rationality: Evidence and theories of preferential choice. *Journal of Economic Literature*, 44(3), 631–661.
- Russo, J. E. (2015). The predecisional distortion of information. In E. A. Wilhelms & V. F. Reyna (Eds.), *Neuroeconomics, judgment, and decision making* (pp. 91–110). Psychology Press.
- Russo, J. E., Carlson, K. A., Meloy, M. G., & Yong, K. (2008). The goal of consistency as a cause of information distortion. *Journal of Experimental Psychology: General*, 137(3), 456–470.
- Russo, J. E., & Chaxel, A.-S. (2010). How persuasive messages can influence behavior without awareness. *Journal of Consumer Psychology*, 20(3), 338–342.
- Russo, J. E., Medvec, V. H., & Meloy, M. E. (1996). The distortion of information during decisions. *Organizational Behavior and Human Decision Processes*, 66, 102–110.
- Salti, M., El Karoui, I., Maillet, M., & Naccache, L. (2014). Cognitive dissonance resolution is related to episodic memory. *PLOS ONE*, 9(9), Article e108579.
- Schonberg, T., Bakkour, A., Hover, A. M., Mumford, J. A., Nagar, L., Perez, J., & Poldrack, R. A. (2014). Changing value through cued approach: An automatic mechanism of behavior change. *Nature Neuroscience*, 17(4), 625–630.
- Schonberg, T., & Katz, L. N. (2020). A neural pathway for nonreinforced preference change. *Trends in Cognitive Sciences*, 24(7), 504–514.
- Sharot, T., Fleming, S. M., Yu, X., Koster, R., & Dolan, R. J. (2012). Is choice-induced preference change long lasting? *Psychological Science*, 23(10), 1123–1129.
- Shohamy, D., & Adcock, R. A. (2010). Dopamine and adaptive memory. *Trends in Cognitive Sciences*, 14(10), 464–472.
- Simon, D. (2012). More problems with criminal trials: The limited effectiveness of legal mechanisms. *Law and Contemporary Problems*, 75(2), 167–209.
- Simon, D., & Holyoak, K. J. (2002). Structural dynamics of cognition: From consistency theories to constraint satisfaction. *Personality and Social Psychology Review*, 6, 283–294.
- Simon, D., Krawczyk, D. C., Bleicher, A., & Holyoak, K. J. (2008). The transience of constructed preferences. *Journal of Behavioral Decision Making*, 21, 1–14.
- Simon, D., Krawczyk, D. C., & Holyoak, K. J. (2004). Construction of preferences by constraint satisfaction. *Psychological Science*, 15, 331–336.
- Simon, D., Pham, L. B., Le, Q. A., & Holyoak, K. J. (2001). The emergence of coherence over the course of decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 1250–1260.
- Simon, D., Snow, C. J., & Read, S. J. (2004). The redux of cognitive consistency theories: Evidence judgments by constraint satisfaction. *Journal of Personality and Social Psychology*, 86(6), 814.

- Simon, D., & Spiller, S. A. (2016). The elasticity of preferences. *Psychological Science*, 27(12), 1588–1599.
- Simon, D., Stenstrom, D. M., & Read, S. J. (2015). The coherence effect: Blending hot and cold cognitions. *Journal of Personality and Social Psychology*, 109, 369–394.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3), 665–690.
- Spellman, B. A., Ullman, J. B., & Holyoak, K. J. (1993). A coherence model of cognitive consistency. *Journal of Social Issues*, 4, 147–165.
- Svenson, O. (1992). Differentiation and consolidation theory of human decision making: A frame of reference for the study of pre- and post-decision processes. *Acta Psychologica*, 80, 143–168.
- Tversky, A. (1969). Intransitivity of preferences. *Psychological Review*, 76(1), 31–48.
- Tversky, A., & Simonson, I. (1993). Context-dependent preferences. *Management Science*, 39(10), 1179–1189.
- Voigt, K., Murawski, C., Speer, S., & Bode, S. (2019). Hard decisions shape the neural coding of preferences. *The Journal of Neuroscience*, 39(4), 718–726.
- Warren, C., McGraw, A. P., & Van Boven, L. (2011). Values and preferences: Defining preference construction. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(2), 193–205.

Received August 20, 2020

Revision received February 13, 2021

Accepted February 16, 2021 ■