Decision

Transient Value Refinements During Deliberation Facilitate Choice

Douglas G. Lee and Keith J. Holyoak

Online First Publication, May 11, 2023. https://dx.doi.org/10.1037/dec0000215

AMERICAN PSYCHOLOGICAL ASSOCIATION

Decision

ISSN: 2325-9965

https://doi.org/10.1037/dec0000215

Transient Value Refinements During Deliberation Facilitate Choice

Douglas G. Lee[1] and Keith J. Holyoak[2]

[1] School of Psychological Sciences, Tel Aviv University

[2] Department of Psychology, University of California, Los Angeles

After deciding among options, decision-makers tend to increase their evaluations of the chosen options and decrease their evaluations of the rejected options, resulting in a *spreading of alternatives* (SoA). There has been a long-standing debate as to whether SoA results from postchoice cognitive dissonance reduction or self-consistency augmentation, or whether it is related to processes that are instrumental in reaching a decision (or both). Here, we introduce a novel procedure that measures SoA implicitly during the process of subjective value-based decision-making. During the choice task, participants simultaneously provided value ratings for both options on offer on each trial, but *before* explicitly reporting their choices. The results clearly demonstrate that SoA does not occur only after choices have been reported. Furthermore, SoA seems to be instrumental to the choice process: It is associated with higher choice consistency, higher confidence, and lower response time. The SoA generated during choice deliberation is partially transient in nature, regressing toward baseline shortly afterward, but lingering at an intermediate level.

Keywords: decision-making, spreading of alternatives, choice-induced preference change, value refinement, confidence

Supplemental materials: https://doi.org/10.1037/dec0000215.supp

That people make decisions based on their personal preferences is undisputable—when deciding among options, we will always try to choose the one we believe to be the best for ourselves under the given circumstances. What is less well understood is the opposite pattern, in which, *after* deciding among options, we tend to increase our preference for the one we chose and decrease our preference for those we rejected. This phenomenon of *choice-induced preference change* (a variety of what has been termed *coherence shift*) has been studied for more than 60 years, exploiting the *free-choice paradigm* to demonstrate such preference change via the *spreading of alternatives* (SoA). The standard paradigm consists of three successive tasks: prechoice ratings of individual options, choices between pairs of options, and postchoice ratings of individual options (Brehm, 1956). SoA is based on the observation that the difference in the subjective value ratings that people assign to options in a choice pair is typically larger after the choice than before (i.e., the values "spread apart"). This effect is highly robust (e.g., Carlson & Russo, 2001; Chammat et al., 2017; Holyoak & Simon, 1999; Izuma et al., 2010, 2015; Lee & Coricelli, 2020; Lee & Daunizeau, 2020, 2021; Lee & Hare, 2023; Lee & Holyoak, 2021; Russo

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

This work was supported by a Google Faculty Research Award to Keith J. Holyoak.

Douglas G. Lee played lead role in conceptualization, data curation, formal analysis, investigation, methodology, project administration, resources, software, validation, visualization, writing–original draft, and writing–review and editing. Keith J. Holyoak played lead role in supervision and supporting role in conceptualization, formal analysis, investigation, methodology, project administration, writing–original draft, and writing–review and editing.

The raw data and primary analysis code for this study are available in the Open Science Forum at http://doi.org/10.17605/OSF.IO/PQW4X (Lee, 2021).

Correspondence concerning this article should be addressed to Douglas G. Lee, School of Psychological Sciences, Tel Aviv University, Tel Aviv 6997801, Israel. Email: douglas@tauex.tau.ac.il

et al., 1996; Salti et al., 2014; Sharot et al., 2009, 2010, 2012; Voigt et al., 2017, 2019; see Enisman et al., 2021, for a meta-analysis).

SoA has commonly been viewed as a form of cognitive dissonance reduction: When people choose one option over an alternative that they thought they had liked approximately equally well, they sometimes feel a sense of psychological discomfort (dissonance or regret). They then attempt to reduce the dissonance (or avoid the regret) by convincing themselves afterward that they actually preferred the chosen option more than they had previously thought (Festinger, 1957; for a review, see Harmon-Jones & Mills, 2019). But other work over the decades has demonstrated that some form of SoA may take place *while* the decision process is unfolding (Carlson & Russo, 2001; DeKay, 2015; DeKay et al., 2012; Holyoak & Simon, 1999; Russo et al., 1996, 1998, 2008). Such effects have been reported in medical decisions (Kostopoulou et al., 2012; Nurek et al., 2014); legal decisions (Carlson & Russo, 2001; Engel & Glöckner, 2013; Holyoak & Simon, 1999; Simon et al., 2001; Simon, Snow, et al., 2004); consumer decisions (Carlson et al., 2006; Russo et al., 1996, 1998, 2008); as well as employment, professional, entrepreneurial, and educational decisions (Miller et al., 2013; Russo, 2015).

More recently, additional evidence has amassed suggesting that SoA is not solely the result of postchoice dissonance reduction, but rather takes place *during* choice deliberation (not only after) and is instrumental to the choice itself (Voigt et al., 2019). Lee and Daunizeau (2021) proposed a model that explains SoA based on three assumptions: (a) prechoice value estimates are imprecise, (b) choices between similarly valued options encourage refinement of value estimates, and (c) postchoice value estimates retain the higher precision and accuracy achieved during choice deliberation. This model assumes that changes that take place during choice deliberation will endure and thus be observable in the subsequent rating task. However, it is not obvious that this should be the case. In fact, most prior studies have not investigated the temporal dynamics of preference change, relying instead on initial and final value measurements to demonstrate that some degree of change did indeed occur at some point. A notable exception is the work on *information distortion* reported by Russo and colleagues (Carlson & Russo, 2001; Russo et al., 1996). That work showed that as preferences are being formed over the course of deliberation, the value provided by *newly presented* information is distorted in favor of either preexisting or initially formed preferences. Thus, a specific form of preference change was observed, wherein evolving value estimates were pulled in the direction of the initial (or current) preference.

Other work has shown that statistical artifacts related to how SoA is measured can account for some of the observed effects, even when no actual preference change has occurred (Chen & Risen, 2010). The basic idea is that, under the assumptions that ratings are noisy estimates of true preferences and choices reflect true preferences, postchoice ratings of choice option pairs will be (on average) farther apart relative to the corresponding prechoice ratings. It has been demonstrated via simulation that the statistical observation holds when the difference in prechoice ratings is sufficiently small (Alós-Ferrer & Shi, 2015; Izuma & Murayama, 2013; Lee & Pezzulo, 2022; see also Supplemental Material). However, many studies have since shown that SoA still occurs even after controlling for such statistical artifacts.

In the present study, we seek to demonstrate that the preference change occurs *during* choice deliberation using a novel design that can provide strong evidence in support of that claim. In brief, we solicit value estimates for the choice options at the same time as the choices themselves, which makes it possible to determine whether SoA has already occurred by the time the choice response is reported.[1] This experimental design is similar in spirit to that of the information distortion studies, except that in that paradigm, new information is repeatedly and explicitly presented for the decision-maker to consider. In other words, that paradigm only measures predecisional changes in preferences that arise in direct response to the introduction of additional explicit information. In our paradigm, no new information is ever presented, so

[1] If so, SoA might be instrumental to the choice, in that it allows for choices to be made more quickly and with higher confidence (as reported in previous studies). But SoA may be more than the product of a bias, driven by the implicit goal of conserving effort or maximizing confidence. If SoA simply reflected a bias, it would not be instrumental to the choice (i.e., it would not help the decision-maker choose the more valuable option).

that any new information that might possibly be considered must arise from the decision-maker's own mind.

Our paradigm also allows us to test a second hypothesis: That choice-induced preference change is at least partially transient in nature. It is possible that apparent differences in the subjective value of choice options might be especially large during deliberation, enabling decision-makers to choose more easily (i.e., without deliberating excessively). The immediate SoA might have already partially faded by the time that final value judgments are elicited. There is some evidence that SoA might diminish over time, on the order of 10 days (Simon & Spiller, 2016) or only 15 min when distractor tasks are interleaved (Simon et al., 2008), but our paradigm allows us to examine whether SoA begins to dissipate between the time of choice and the final rating task that immediately follows. To preview our results, this is exactly what we found: SoA during choice deliberation was associated with increased choice consistency and confidence and reduced response time; SoA then dissipated after the choice, with value estimates regressing (though only partially) back toward their initial levels. In addition to testing our hypotheses, we also check whether the data could be explained as a statistical artifact; we show that it cannot.

Material and Method

We examined shifts in subjective value triggered by choices between pairs of snack foods. The design included three experimental conditions: standard choice task, implicit choice task using simultaneous ratings, and standard plus implicit choice task. The purpose of this study was to observe choice-induced preference changes as they occurred during choice deliberation, rather than only later during the subsequent rating task. This design also allowed us to investigate the transient nature of choice-induced preference changes, by recording them as they occurred (during the choice task itself) and then measuring their residual effects in the subsequent rating task. The key experimental condition was thus the implicit choice task. We included the standard choice task to serve as a baseline, as results under that condition have been reported many times before. We included the standard plus implicit choice task to verify that including rating scales during the choice, task does not interfere with the usual pattern of results with respect to the choices and the postchoice ratings.

Participants

A total of 178 people participated (90 female; age: $M = 42$ years, $SD = 8$, range = 24–55 years), split evenly across three experimental conditions (*standard choice*: $n = 60$, 28 female, age: $M = 41$ years, $SD = 8$, range = 26–55; *implicit choice*: $n = 60$, 30 female, age: $M = 41$ years, $SD = 8$, range = 24–55; *standard + implicit choice*: $n = 58$, 32 female, age: $M = 43$ years, $SD = 9$, range = 27–55). This sample size (per condition) was chosen to be comparable to that used in previous studies based on a similar paradigm. All participants were recruited using the Amazon Mechanical Turk (MTurk; https://mturk.com). All were classified as "masters" by MTurk. They were all residents of the United States or Canada, and all were native English speakers. Each participant received a payment of $7 as compensation for approximately 45 min 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 (https://gorilla.sc). The experimental stimuli were the same set of 100 digital images used in a previous study (Lee & Holyoak, 2021), each representing a distinct snack food.

Design and Procedure

The experiment consisted of four phases: preexposure, initial value rating, choice, and final value rating (see Figure 1, for an illustration of the experimental design). No time limits were imposed for any of the constituent tasks or for the overall experiment. After providing informed consent, the participants were presented with the following instructions:

> In this study, you will be asked to provide simple information regarding your preferences for consuming different types of snack foods. Specifically, for each of a variety of snacks, you will tell us how much you feel

Figure 1
Experimental Design

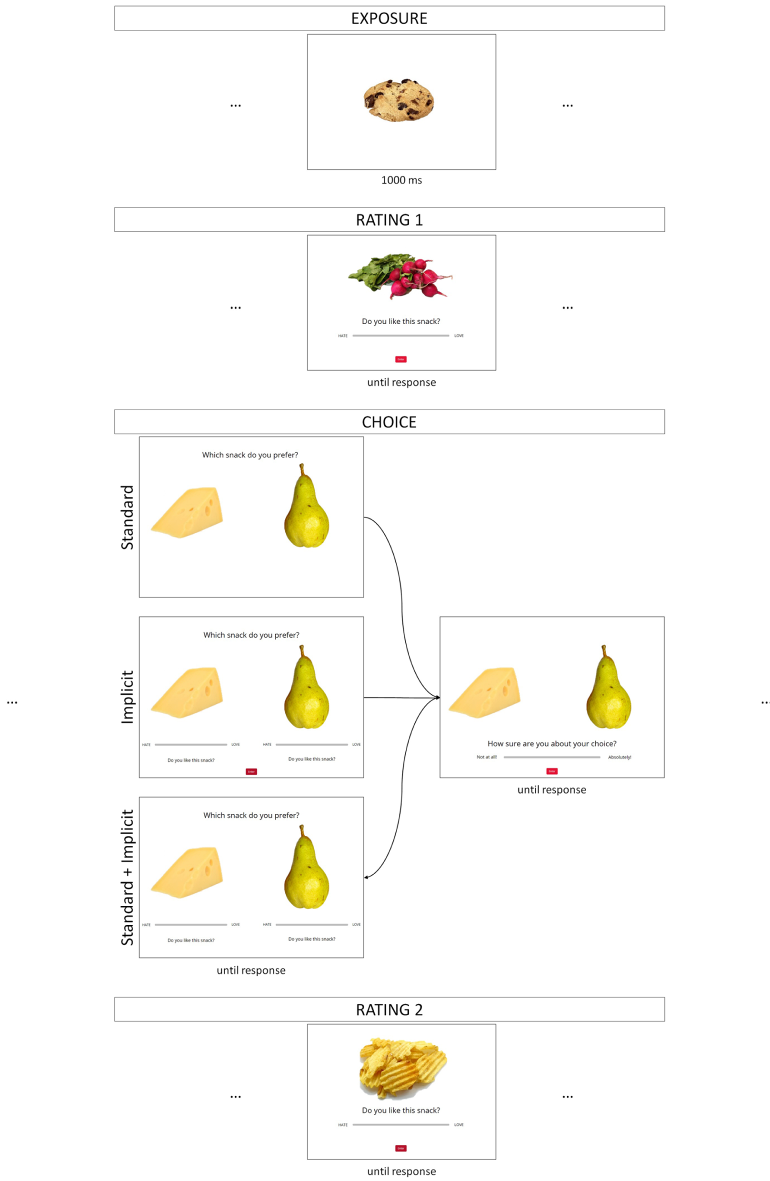

Note. In the exposure task, participants passively viewed each option for 1 s. In the first and second rating tasks, participants reported how much they liked each option at their own pace. In the choice task, participants first reported their preferred option within each pair and then their confidence about each reported preference. In the standard condition, participants clicked on the image of their preferred option. In the implicit condition, participants entered individual likability ratings on separate slider scales and then clicked an "enter" button, so that the ratings themselves served as an implicit choice. In the standard + implicit condition, participants entered individual likability ratings on separate slider scales before clicking on the image of their preferred option. In all conditions, participants entered their confidence in each choice immediately after making the choice. Images from Pixabay (https://www.pixabay.com/). See the online article for the color version of this figure.

> that you would enjoy eating it on a frequent basis. You will also be asked to choose your preferred item from different pairs, as well as to estimate how confident you are about each choice.

The instructions that were presented to participants at the beginning of each task are provided in Supplemental Material.

In the preexposure phase, participants simply observed as all individual items were displayed in a random sequence for 1,000 ms each (with no intertrial interval). The purpose of the preexposure 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 value across the item set. Just prior to the onset of the preexposure phase, participants were provided with instructions related to the task.

In the initial value rating task (for all three conditions), all stimuli were displayed on the screen, one at a time, in a sequence randomized across participants. Before the onset of this task, the participants were provided with instructions related to the task. 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 this rating of overall value, participants responded to the question, "Do you like this snack?" using a horizontal slider scale. The instructions for this task encouraged participants to think carefully while assessing the overall subjective quality of each option, by asking them to imagine that the choice was for frequent consumption, rather than a "one-off" snack. The leftmost end of the scale was labeled "HATE," and the rightmost end was labeled "LOVE." The scale appeared to participants to be continuous, and the response 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 choice task was then administered, with different variants for each condition. For this task, 50 pairs of snacks were displayed on the screen, one pair at a time, in a sequence randomized across participants. The pairings of snacks for each choice trial were identical to those used in a previous study (Lee & Holyoak, 2021). These pairings had been 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 as measured in a previous study (Lee & Coricelli, 2020). To clarify, all ratings (and thus all measures of choice difficulty) in the present study were specific to each participant; only the standardized set of choice pairs (i.e., the stimulus set) was determined based on ratings from the previous study. Each individual item occurred within 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 snack foods appeared on the screen, one left and one right of center.

Participants were randomly assigned to one of the three conditions for this task, counterbalanced across conditions. In all three conditions, the participants responded to the question, "Which snack do you prefer?" Before beginning the choice task, participants received instructions related to the specific condition to which they were assigned. In the standard choice condition, participants responded by clicking on the image of their preferred item (as in Lee & Holyoak, 2021). In implicit choice, participants responded by providing ratings for each individual item on independent slider scales before clicking "enter" to continue. Each slider scale was located under its respective snack food and was identical in format to the one used in the isolated rating task (other than being half the physical length). In standard + implicit choice, participants first entered simultaneous individual ratings (as in implicit choice), but then finalized their choice by clicking on the option they preferred (as in standard choice), rather than the "enter" button. All 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 in all three conditions made final ratings of overall value, exactly as for the initial ratings, except that the stimuli were presented in distinct random orders. Note that this procedure (randomizing the order of individual items) serves to dissociate the final ratings from the context of the initial ratings (reducing any tendency to try to recall the initial ratings when making final ratings). 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 currently evaluated them:

> NOTE: Please respond according to how you feel at the present moment for any given item, regardless of how you might have felt at any other time. Do not try to match what you might have responded in other ratings, that will not matter to us and it could actually spoil the data.

Exclusion Criteria

Due to the difficulty of experimental control in online experiments, we anticipated that participants might not pay full attention to the task at hand on every trial. We therefore excluded from analysis all trials with an outlier response time (RT). Specifically, we calculated the median and median average deviation of RT across all trials and all participants within each condition. We excluded any trial whose RT was more than three times the median average deviation away from the median (within each condition). This resulted in the exclusion of 139, 131, and 112 trials in the standard, standard + implicit, and implicit conditions, respectively.

Statistical Analyses

All regression results reported below were calculated by mixed-effects linear regression using the *fitlme* Matlab function, with random-effects slope and intercept terms for participants and fixed effects for all variables of interest. For all other (nonregression) results, the reported *p* values are based on standard two-sided *t* tests. To assist with both readability and interpretation, we coded all variables such that the left option (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 (left option minus right option). Because our design involves two standard rating phases (pre- and postchoice), we distinguish the rating difference from these phases by labeling them dV1 (prechoice) and dV2 (postchoice). The implicit choice and standard + implicit choice conditions also involved an additional rating phase (intrachoice). We label the rating difference from this phase as dV^*.

The primary focus of our analyses is on choice-induced preference change, which results in the SoA between the initial and final subjective value ratings. The choice defines the winning option, and SoA is defined in terms of changes that relatively favor the winner. Specifically, SoA is defined as the change in rating for the chosen option (from initial to final rating) minus the change in rating for the unchosen option. Again, because our experimental design involves multiple rating phases, we distinguish the SoA that occurs between the initial rating and the choice phases (which we label SoA^*) from the residual SoA that occurs between the choice and the final rating phases (which we label SoA^R). The traditional measure (between the initial and final rating phases) retains the label SoA, such that $SoA = SoA^* + SoA^R$.

Apart from SoA, we also examine choice consistency (defined as a choice in favor of the option that was rated with a higher value), RT (measured from the presentation of the choice options until the response), and confidence (for each choice). These variables always relate to the choice phase of the experiment and are independent from any of the rating phases. We assess the relationships between choice difficulty and each of the behavioral choice variables (consistency, RT, SoA, and confidence). We define difficulty according to the difference in initial value ratings between the options (dV1), where higher values of |dV1| imply lower choice difficulty. We then separately regress consistency (logistic) on dV1 and RT, SoA, and confidence on |dV1|, testing the beta weights for significance. We rescaled all relevant variables from a (0,100) scale to a (0,1) scale before adding them to the design matrices.

To test whether the ratings obtained in the final round of the experiment were better predictors of the choice variables, we included dV2 as a coregressor in all models (apart from the SoA model, because dV1 and dV2 by definition have the opposite relationship with SoA by definition). Because dV1 and dV2 are highly correlated, we first regressed dV2 on dV1 and entered the residuals as regressors in our main models of interest. To test whether the ratings obtained at the time of choice (in the implicit and standard + implicit conditions) were better predictors of the choice variables, we included dV^* as a coregressor in all models (apart from SoA). Because dV1, dV2, and dV^* are highly correlated, we first regressed dV^* on dV1 and dV2 and entered the residuals as regressors in our main models.

For certain analyses, we categorize decisions as either "difficult" or "easy." In line with previous studies, we define a difficult choice as one in which dV1 is less than 10 points (on the 100-point rating scale).

Ruling Out the Statistical Artifact Explanation

Chen and Risen (2010) pointed out that a statistical artifact in the way in which the SoA is calculated can produce an apparent effect even when no true preference change takes place. Specifically, if the value ratings that participants report are assumed to be noisy measures of the true underlying subjective values, and the choices that participants report are assumed to align with their true preferences, then regression to the mean can sometimes cause the ratings of choice pair options to spread apart from pre- to postchoice. Chen and Risen provided a mathematical demonstration that positive SoA is always predicted even when preferences are stable, and other authors have further explained the finding conceptually and via computer simulation (Alós-Ferrer & Shi, 2015; Izuma & Murayama, 2013).

Lee and Pezzulo (2022) recently proposed a simple method to test whether observed SoA is caused by a genuine cognitive phenomenon rather than simply being a statistical artifact. Their method involves simulating the free-choice paradigm under different assumed models: noise only, cognitive dissonance, or value refinement (among other variations). Although efforts to distinguish the different cognitive models were not fully conclusive, the method provided clear evidence that the SoA observed in several previous studies could not be explained by the statistical artifact model. Here, we rely on this method to reject the possibility that the SoA observed in our results might be nothing more than a statistical artifact. We summarize the procedure we followed in the Supplemental Material.

Results

We first examined the fundamental dependent variables (choice consistency, response time, choice confidence) in our three experimental choice conditions. Across participants, the mean consistency was 83%, 84%, and 84% in the standard, standard + implicit, and implicit conditions, respectively. Across participants, the mean RT was 1.8 s, 5.4 s, and 5.4 s in the standard, standard + implicit, and implicit conditions, respectively. Responses took substantially longer in the standard + implicit and implicit conditions because participants had to enter their responses on the slider scales prior to clicking to confirm their choice. Across participants, the mean confidence was 74%, 77%, and 83% in the standard, standard + implicit, and implicit conditions, respectively. The difference in confidence across all three conditions was significant (all $p < .001$ using standard two-tailed t tests), suggesting that explicitly evaluating the options at the time of choice might have helped participants feel more sure about their choices.

As a check that participants in the standard + implicit choice condition performed the task as expected, we verified that the chosen option (indicated by explicit selection) matched the higher rated option (during the choice task, indicated by separate slider scales for each option). Indeed, across participants, 91% of trials conformed to this pattern. The mismatch on the remaining trials was likely due to imprecision in the use of the rating scales (participants were not aware of the numerical values associated with points on the scales and had to rely on visual estimation to tune their ratings). The average prechoice value difference (dV1) on such mismatch trials was only 11 points on the 100-point scale; this contrasts with an average dV1 of 26 points on trials in which the explicit and implicit choices matched.

The preference change phenomenon could be such that SoA is always in the direction of initial preferences (in line with previous work on information distortion and coherence shifts; Carlson & Russo, 2001; Holyoak & Simon, 1999; Russo et al., 1996, 2008). Alternatively, SoA could be nondirectional, meaning that it sometimes causes preference reversals (i.e., the chosen option is not the same as the option that was initially rated higher). One might argue that the primary benefit of SoA occurs in those cases where it is instrumental to be inconsistent: The possibility that one might change one's mind about which option one prefers (after careful consideration) is perhaps the most critical aspect of deliberation. Otherwise, there would never be any need to deliberate: One would simply choose based on one's initial preference (however uncertain it might be). Accordingly, apparent changes of mind (i.e., choices inconsistent with prechoice ratings but consistent with postchoice rating) should be more common than other types of inconsistency (i.e., choice consistent with pre- but not postchoice ratings—"rating error"—or consistent ratings with an inconsistent choice—"choice error"). In each experimental condition, changes of mind (CoM) were significantly more common than either type of error (standard: mean CoM = 0.09, mean rating errors [err_r] = 0.06, mean choice errors [err_c] = 0.07; standard + implicit: mean CoM = 0.09, mean err_r = 0.06, mean err_c = 0.06; implicit: mean CoM = 0.09, mean err_r = 0.06, mean err_c = 0.06). These differences are even more substantial when examining only difficult choices (standard: mean CoM = 0.21, mean err_r = 0.10, mean err_c = 0.10; standard + implicit: mean CoM = 0.21, mean err_r = 0.10, mean err_c = 0.12; implicit: mean CoM = 0.18, mean err_r = 0.13, mean err_c = 0.10; see Table 1).

Choice Induces Preference Change

In accord with previous work, we observed a reliable SoA[2] across all choice trials (cross-participant mean of the within-participant mean; standard: SoA = 2.1; standard + implicit: SoA = 2.1; implicit: SoA = 2.2). These values of SoA are comparable to the magnitude reported in previous studies (Izuma et al., 2010; Lee & Daunizeau, 2020; Lee & Holyoak, 2021; Voigt et al., 2019). The effect size did not differ across conditions (all $p > .727$), confirming that our method of elicitation did not alter the processes that give rise to SoA. We then assessed the relationship between choice difficulty and SoA by regressing SoA on |dV1|. As in previous studies (Lee & Coricelli, 2020; Lee & Daunizeau, 2020, 2021; Lee & Hare, 2023; Lee & Holyoak, 2021), |dV1| had a reliable negative relationship with SoA in every experimental condition (Table 2).

Choice-Induced Changes in Preference Are Instrumental to Decisions

Previous research has suggested that decision-makers refine their value estimates for choice options during deliberation, prior to committing to the choice (Lee & Daunizeau, 2020, 2021). It has been shown that the predictive effect of dV on consistency is larger when dV is calculated using postchoice ratings (i.e., dV2) rather than prechoice ratings (Lee & Coricelli, 2020; Lee & Daunizeau, 2020, 2021; Lee & Hare, 2023; Lee & Holyoak, 2021; similar results were also reported in Simon, Krawczyk, et al., 2004; Simon & Spiller, 2016). The present study replicates this finding. Specifically, we first regressed consistency on dV1 and dV2 residuals (see Material and Method section). For the standard + implicit and implicit conditions, we also included dV^* residuals (see Material and Method section) as a regressor. If intrachoice value difference (dV^*) is meaningful, in the sense that SoA^* (the difference between dV1 and dV^*) is instrumental to the choice, we would expect dV^* to have a predictive effect on choice consistency even after controlling for the predictive effect of pre- and postchoice value difference. In other words, we would expect dV^* to be a better predictor of choice behavior than dV1, and not worse of a predictor than dV2 (which would be the case if preferences changed only after the choice, as postulated by traditional—e.g., cognitive dissonance, self-consistency—theories). Unsurprisingly, we found that dV1 has a reliable positive relationship with consistency (Figure 2 and Supplemental Table S1). Critically, in all conditions, we found that dV2 has an incremental predictive effect beyond the effect of dV1 and dV^* has an incremental predictive effect beyond the effect of dV2 (Figure 2 and Supplemental Table S1).

Previous studies have shown that SoA is positively associated with choice confidence (Lee & Daunizeau, 2020, 2021; Lee & Hare, 2023; Lee & Holyoak, 2021). The impact of SoA (if it occurs during deliberation) is to effectively make the choice easier prior to entering a response. In brief, a lower prechoice value difference (dV1) will make the choice relatively difficult and thus encourage deliberation before responding. Deliberation tends to generate SoA (essentially an increment in dV1), which in turn increases confidence that one option is better than the other. We replicated these findings. Specifically, we separately regressed RT and confidence on |dV1| and SoA. For the standard + implicit and implicit conditions, we separated SoA into regressors for each of its components: SoA^* and SoA^R. The presumed

[2] For consistent choices, SoA = dV2 − dV1; for changes of mind, SoA = dV1 − dV2.

Table 1
Comparisons of Changes of Mind Versus Rating and Choice Errors, Across Experimental Conditions

Dependent variable	*M*	CI	*t*	*df*	*p*
Standard					
CoM - err_r	0.032	[0.0163, 0.0478]	4.071	59	<.001
CoM - err_c	0.021	[0.0050, 0.0364]	2.638	59	.011
CoM - err_r (difficult only)	0.103	[0.0678, 0.1388]	5.827	59	<.001
CoM - err_c (difficult only)	0.103	[0.0698, 0.1371]	6.147	59	<.001
Standard + Implicit					
CoM - err_r	0.035	[0.0192, 0.0503]	4.463	57	<.001
CoM - err_c	0.029	[0.0145, 0.0431]	4.027	57	<.001
CoM - err_r (difficult only)	0.113	[0.0736, 0.1527]	5.729	57	<.001
CoM - err_c (difficult only)	0.096	[0.0574, 0.1337]	5.014	57	<.001
Implicit					
CoM - err_r	0.028	[0.0125, 0.0437]	3.608	59	<.001
CoM - err_c	0.023	[0.0057, 0.0411]	2.643	59	.011
CoM - err_r (difficult only)	0.052	[0.0135, 0.0905]	2.703	59	.009
CoM - err_c (difficult only)	0.080	[0.0400, 0.1205]	3.988	59	<.001

Note. *M* = mean; CI = confidence interval; CoM = change of mind rate; err_r = rating error rate; err_c = choice error rate.

increase in option discriminability that arises during choice deliberation is captured by SoA*, so we expect SoA* to have a predictive effect on RT and confidence. With respect to SoA^R, it is less clear what to expect. It might be that SoA^R has no predictive effect on choice, because it represents changes in value estimates that occur after the choices are reported. Or, SoA^R might have a negative predictive effect, if it was merely caused by noise unrelated to choice. Finally, SoA^R might have a positive predictive effect, if it reflects a sort of correction against exaggerated coherence shifts (i.e., SoA* resulting in dV* greater than the true subjective value difference). Unsurprisingly, we found that |dV1| has a reliable positive relationship with confidence and a reliable negative relationship with RT (Figure 3 and Supplemental Tables S2 and S3). Critically, for both dependent variables in all conditions, we found that SoA* has an incremental predictive effect beyond the effect of |dV1| (Figure 3 and Supplemental Tables S2 and S3). Interestingly, for both dependent variables in both relevant conditions, we also found that SoA^R has an incremental predictive effect beyond the effect of |dV1| and SoA* (Figure 3 and Supplemental Tables S2 and S3). All beta weights for SoA resembled those for |dV1| (though with smaller magnitudes), demonstrating that those variables serve a similar role in determining choice behavior.

Table 2
Regression Coefficients Relating Choice Ease and SoA, Across Experimental Conditions

Regression coefficients	*M*	*SE*	*t*	*df*	*p*
Standard					
β |dV1| for SoA	−0.13	0.02	−7.51	2,859	<.001
Standard + Implicit					
β |dV1| for SoA	−0.21	0.02	−10.19	2,767	<.001
Implicit					
β |dV1| for SoA	−0.21	0.02	−10.72	2,886	<.001

Note. *M* = mean; *SE* = standard error; dV1 = prechoice value difference; SoA = spreading of alternatives.

It is our contention that SoA occurs during choice deliberation and is instrumental to the developing choice, with residual effects that persist beyond the conclusion of the choice. Accordingly, dV2 should provide a more predictive power than dV1, as the postchoice ratings should include a more decision-relevant information than the prechoice ratings. If postchoice ratings are richer in decision-relevant information compared to prechoice ratings, the proportion of the variation in all measures of choice behavior (consistency, RT, and confidence) that is predictable from dV should be higher when using ratings obtained after all choices have been made. We thus ran a series of separate regressions of all choice variables on dV1 and

Figure 2
dV2 and dV Incrementally Explain Choice Beyond dV1*

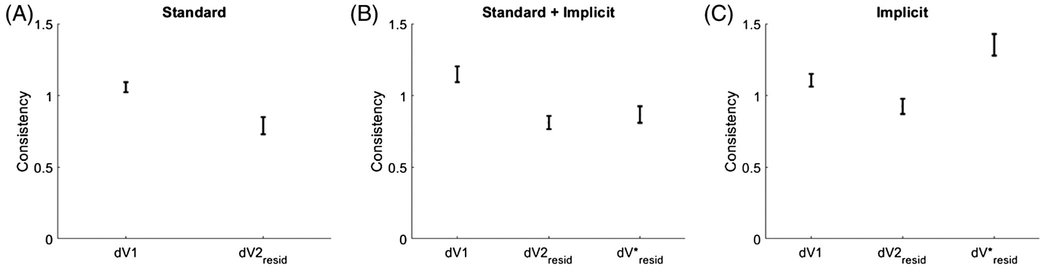

Note. Predictive effect of dV1, dV2 residuals, and dV* residuals on choice, across conditions (columns; recall that dV* does not apply in the standard condition). For dV2, dV1 was first regressed out. For dV*, dV1 and dV2 were first regressed out. These coefficients thus represent the added predictive effect of dV2 above and beyond dV1 and the added predictive effect of dV* above and beyond dV1 and dV2. Error bars represent 95% confidence intervals. dV = value difference.

dV2 and examined the *r*-squared terms. (For the regressions of RT and confidence, we first took the absolute value of dV.) Across all variables and all conditions, the adjusted *r*-squared was larger for the regressions based on dV2 (Table 3).

Choice-Induced Changes in Preference Are Partially Transient

The central hypothesis that we tested in this study is that choice-induced preference change occurs *during* choice deliberation (and not only after the choice has been made). Our methodological approach was to measure preference ratings and solicit choices simultaneously. This design provided us with a novel variable for the standard + implicit and implicit conditions: SoA *during* choice deliberation (SoA*), which is important for the emerging choice. We used the simultaneous value ratings obtained during the choice task in the standard + implicit and implicit conditions to compute intermediate (i.e., at the time of choice) measures of dV (which we label dV*). The measures of dV* are entered into the SoA* calculation in exactly the same way as the postchoice value difference (dV2) is entered into the regular SoA calculation. We observed a reliable SoA* across all choice trials (standard + implicit choice: cross-participant mean of within-participant mean SoA* = 4.0, $p < .001$; implicit choice: cross-participant mean of within-participant mean SoA* = 6.5, $p < .001$).

Note that the relatively large magnitude of SoA* compared to the final SoA (in both standard + implicit and implicit conditions) demonstrates that there was a *negative* SoA between the choice task and the final rating task. This reduction in SoA indicates that choice-induced preference changes are partially transient in nature: The perceived values of the individual options regress back toward their initial (perhaps less precise) values once they are removed from the choice (or other comparative) context. The SoA observed in standard paradigms might therefore be residual in nature, with traces of the information that enabled the choice between options remaining more easily accessible during subsequent evaluations. One possible interpretation is that revisions of value difference estimates are exaggerated during choice deliberation in order to facilitate the decision, but then the exaggeration deflates after the choice is made (see Figure 4 for an illustrative example).

To quantify the transient nature, we separated the final SoA into its component parts: SoA* and the residual change SoA^R. Thus defined, the relationship between SoA* and SoA^R serves as an indication of the extent to which choice-induced preference changes are transient and only present during choice deliberation (with value estimates returning to their prior states thereafter). A linear relationship of −1 would imply that choice-induced preference changes were not long-lasting and had no impact on subsequent evaluation tasks (on average across trials). A linear relationship of 0 would imply that choice-induced preference changes were robust, maintaining their full impact on subsequent tasks (on average across trials). We found that the cross-participant mean correlation between SoA* and SoA^R (a measure of SoA transience) was −0.50 ($p < .001$) in the standard + implicit condition and −0.43 ($p < .001$) in the implicit condition (see Figure 5).

Figure 3
SoA Incrementally Explains RT and Confidence Beyond |dV1|

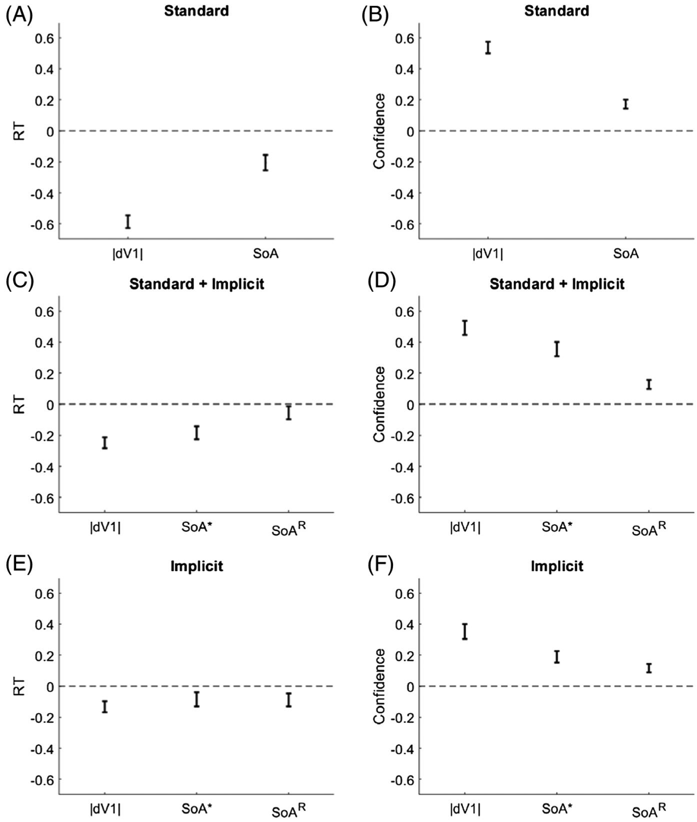

Note. Predictive effect of |dV1|, SoA*, and SoAR on RT (left plots) and confidence (right plots), across conditions (rows; recall that the separation of SoA into SoA* + SoAR does not apply in the Standard condition). Error bars represent 95% confidence intervals. dV1 = prechoice value difference; SoA = spreading of alternatives; RT = response time.

Discussion

The SoA between choice options is a robust phenomenon that has been reported numerous times over the decades. Early accounts of this phenomenon held that people change their evaluations *after* making their choice, either to relieve the unpleasant feeling associated with cognitive dissonance (i.e., if they disliked some aspects of the option they chose, or liked some aspects of the option they rejected; Festinger, 1957), or to maintain a sense of self-consistency (i.e., they chose one option over the other, so they must like it better; Bem, 1967, 1972). More recent accounts of the SoA phenomenon hold that people change their evaluations *during* the choice process (i.e., while deliberating about which option to choose), and that such changes enable people to decide more accurately, more quickly, and more confidently (Lee & Coricelli, 2020; Lee & Daunizeau, 2020, 2021; Lee & Hare, 2023; Lee & Holyoak, 2021). A separate body of literature has shown that people adjust their evaluations of choice options in the direction of the emerging choice, as if the evidence in favor of one option gathers momentum and causes subsequent information to gravitate in the same direction (Glöckner et al., 2010; Holyoak & Simon, 1999; Simon et al., 2001; Simon, Krawczyk, et al., 2004; Simon, Snow, et al., 2004). Key among previous studies are those involving information distortion, which have

Table 3
Postchoice Ratings Explain Choice Data Better Than Prechoice Ratings

Dependent variable	R^2: Explanatory variable = dV1	R^2: Explanatory variable = dV2
Standard		
Consistency	0.43	0.48
RT	0.36	0.36
Confidence	0.36	0.39
Implicit		
Consistency	0.47	0.52
RT	0.36	0.36
Confidence	0.36	0.41
Standard + Implicit		
Consistency	0.44	0.50
RT	0.43	0.44
Confidence	0.45	0.49

Note. For all dependent variables across all conditions, a larger percentage of variance in the data was explained by models based on postchoice ratings (dV2) versus prechoice ratings (dV1). dV = value difference; RT = response time.

demonstrated that decision-makers are biased such that the relative value of newly considered information regarding choice options is distorted so as to conform with the value of previously considered information (Carlson & Russo, 2001; DeKay, 2015; DeKay et al., 2012; Russo et al., 1996, 1998, 2008). Other studies have reported similar effects, where intermittent choices seemed to establish a bias in subsequent decisions (between new options of similar form as the initial options) such that evidence consistent with the initial choice was selectively enhanced (Talluri et al., 2018). This apparent bias might help the decision system work more efficiently, with attention selectively allocated to those value-related signals considered to be most relevant for the current choice context (Schonberg & Katz, 2020). Previous studies have considered various ways in which intradecision information processing might be biased (see Brownstein, 2003, for a review), which could potentially include the emergence of SoA.

The present study introduces a novel experimental design for the free-choice paradigm. This design has a similar conceptual objective as that of the coherence shift and information distortion studies. However, in those paradigms, new information is repeatedly and explicitly presented for the decision-maker to consider, and thus, the predecisional preference changes that they measure only arise in direct response to the introduction of additional explicit information. Our paradigm differs in that no new information is ever presented, so that any new information that might possibly be considered must arise from the decision-maker's own mind. Moreover, previous paradigms were designed to overtly encourage the construction of preferences by presenting new information and asking for revised evaluations. Our paradigm, in contrast, avoids any such influence on the strategies used by participants to make their decisions. Another key distinction is that previous paradigms have measured preference change by examining changes in evaluations of *information* or *attributes*, whereas we directly examine changes in evaluations of the choice options themselves.

Figure 4
Value Difference Estimates Are Exaggerated During Choice Deliberation

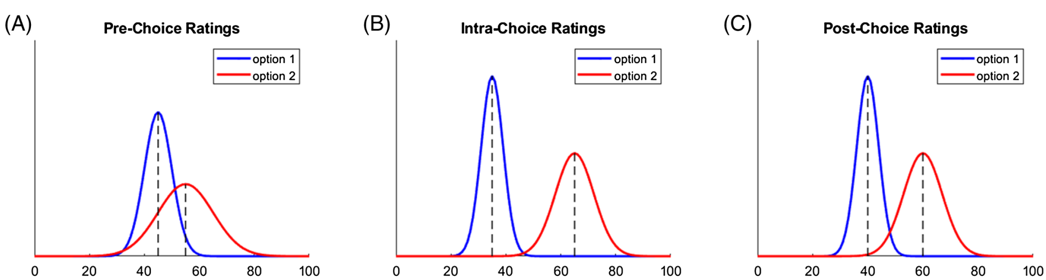

Note. An illustration of how choice-induced preference change might arise. (A) Prechoice ratings show an estimated value difference (dV1) of 10. (B) Intrachoice ratings that inform the decision itself show an exaggerated value difference (dV*) of 30. (C) Postchoice ratings show an estimated value difference (dV2) of 20. In this extreme (for illustrative purposes) example, SoA* is 20 and the final SoA is 10, with a transiency of 50% as in the experimental data. dV = value difference; SoA = spreading of alternatives. See the online article for the color version of this figure.

Figure 5
Transience of Spreading of Alternatives (SoA)

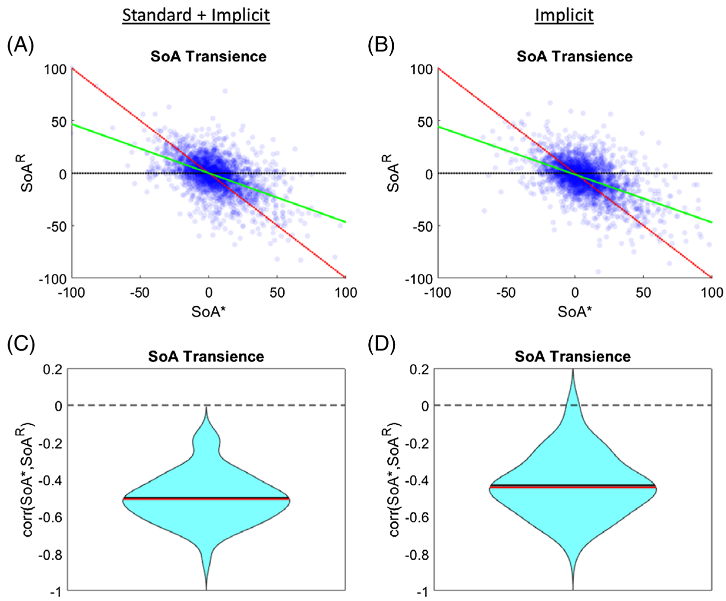

Note. Across all participants, choice pairs that exhibited a larger SoA* (during choice deliberation) exhibited a smaller SoA^R (between choice and final rating. Left panels: Green lines show a linear relationship across all trials and participants, black dotted lines indicate null transiency, red dotted lines indicate full transiency). This pattern resulted in a strong negative correlation between SoA* and SoA^R (right panels). Violin plots represent cross-participant distributions of correlation coefficients; black lines represent cross-participant mean values; and red lines represent cross-participant median values. See the online article for the color version of this figure.

A final difference is that many of the previous paradigms included interim choices or "leanings" solicited before the revised ratings. This means that the revised ratings were in fact postchoice rather than intrachoice. In our paradigm, there are no such interim choices, and thus the ratings we collect during the choice task represent evaluations that influence, not reflect, the choices.

In this study, we provide evidence that SoA occurs before choices are reported, that the SoA is instrumental to the choices during which it occurs, and that the SoA is partially transient in nature. In theory, SoA could occur either before or after the choice is made, or both. We do not provide evidence against the possibility that some postchoice SoA does occur, so we cannot comment on that specifically. However, we do provide strong evidence that a significant amount of SoA occurs *prior* to the choice response. In the implicit choice and standard + implicit choice conditions, participants simultaneously provided value ratings for the options on offer on each trial *before* reporting their choices. These ratings demonstrated a substantial SoA effect in comparison with prechoice ratings, indicating that this SoA did not arise after the choice. Furthermore, the SoA was not a mere statistical artifact (cf. Chen & Risen, 2010), as it was associated with higher choice consistency. It appears that the SoA observed in the present study truly made choices easier (and thus was not random), as SoA was associated with decreased response time and increased choice confidence (Lee & Daunizeau, 2020, 2021). Moreover, postchoice ratings had higher predictive power for choice consistency, RT, and confidence than did prechoice ratings, which would not be the case if the rating changes were mere statistical artifacts. We cannot rule out that some SoA might be caused by noise in the rating process. (Indeed, we believe that to be likely.) However, we can conclude, based on our analyses, that our empirical data cannot be fully explained as an artifact of a statistical effect (see Supplemental Material).

It might be argued that memory may have played some role in causing postchoice ratings to be more consistent with choices than prechoice ratings. One might claim that when rating an option postchoice, a decision-maker simply adjusts the prechoice rating in the direction of the choice outcome (i.e., upward if that option was chosen, downward if it was rejected). However, such a strategy would require that participants recall both Rating1 and Choice for every option in order to produce something qualitatively similar to the observed SoA. This seems extremely implausible, since at the time of Rating2, participants would need to recall 100 previous ratings made up to 40 min prior, and 50 previous choices made up to 20 min prior (in totally random sequences). A similar alternative explanation might suggest that postchoice ratings were more consistent with the choices due to some sort of recency-facilitated recall: At the time of Rating2, the evaluation process that had occurred during Choice would be more recent than that which had occurred during Rating1, causing Rating2 to more closely align with Choice due to a recency effect. However, this cannot be the case, because the timing of Rating1 and Rating2 was symmetrical around Choice. As the temporal distance between either of the rating tasks and the choice tasks was equivalent, any potential recency effect would be equally likely to enhance the consistency of either Rating1 or Rating2 with respect to Choice.

Choices that represent a change of mind (CoM), when the choice and postchoice ratings imply preferences opposite to those implied by prechoice ratings, will necessarily yield observable SoA, as the signs of the relative values of the options will have reversed (i.e., the value of the chosen option will be higher than the value of the rejected option after the choice, but lower before the choice). CoM trials demonstrate that one way in which SoA could arise is when decision-makers realize (after deliberation) that their preferences are actually different than they initially thought. However, CoM trials only represented a small fraction of all trials in the present study, and the SoA effect was still prominent across trials in which there was no CoM.

We note that our data cannot rule out the possibility that participants might first decide (in their minds) which option to choose and then adjust their ratings after the latent choice is finalized but before the actual choice is reported (via the click). Future work could attempt to control for this possibility, perhaps by using process-tracing methods (e.g., eye-tracking) to monitor the within-trial dynamics of attention allocation toward the rating scales. For example, if participants mostly look at the rating scales only at the end of the trial, this might suggest that their internal choices had already been formed, and that they only used the rating scales to conform with their choices. However, if attention alternates between the stimuli and the rating scales throughout the trial, this would suggest that the ratings and choices were developed in tandem. Regardless, the distinction may not be important. It is our belief that ratings and choices, though distinct methods of eliciting external reports of subjective value, are both driven by similar cognitive processes of evaluation. Thus, reporting ratings that align with a just-made choice, or reporting a choice that aligns with just-made ratings, would both reflect the output of the same underlying processes.

With respect to the transient nature of SoA, we showed that SoA* (based on the change in dV from pre- to intrachoice ratings) was substantially greater than the final SoA (based on the change in dV from pre- to postchoice ratings). We cannot rule out the possibility that this effect was partially caused by an abundance of caution on the part of the participants. Knowing that the locations on the slider scales (on which they had to click to report their choices in the implicit choice condition) could not likely be selected with perfect precision, participants may have exaggerated the separation between the ratings in order to ensure that the choice was reported as they intended (thus creating the illusion of SoA). However, we can be sure that this was not the only cause of SoA*, for two main reasons. First, all of the results for the implicit condition matched those of the other conditions, making it unlikely that participants willfully distorted their reported ratings (which would appear as ratings of 0 and 1 on every trial, in the extreme). Second, a significant amount of SoA remained after the final rating task (transiency = 43% in the implicit condition), which would not be expected under this alternative explanation (transiency should be close to 100% in that case). Moreover, there should have been no reason for participants to distort their rating in the standard + implicit choice condition because their choices were reported by clicking on the images of their preferred options (transiency = 50% in this condition, which is even greater than in the implicit condition).

We also demonstrated for the first time that SoA measured via the free-choice paradigm is transient in nature, even with only a single choice task. Judgments of option value spread apart during choice deliberation but then regressed back in the direction of their original values after deliberation ended. This is similar to a previous finding that *attribute* desirability and importance ratings were inflated in the direction of the recently chosen options but dissipated back to baseline after a delay (Simon et al., 2008; Simon & Spiller, 2016). Although we found similar levels of overall SoA (between pre- and postchoice ratings) as did previous studies, the present study is the first to separate SoA into pre- and postchoice SoA. We found that postchoice SoA (between choice and postchoice ratings) was often *negative*, and that the magnitude of negative postchoice SoA was strongly correlated with the magnitude of prechoice SoA. This pattern suggests that although SoA in general is instrumental and based on genuine value-relevant information, it is most pronounced during choice deliberation. As suggested by other work on choice-induced preference change, it seems that the information being considered during choice deliberation is magnified (perhaps by attention) to enable that information to have a maximal impact on the decision, thereby optimizing the efficiency of the system. Then, after the choice, the (attentional) magnification dissipates, and SoA decreases. The residual SoA that lingers beyond the choice may indicate real value estimate refinements, which have been shown to be long-lasting (Sharot et al., 2012).

This interpretation is consistent with a computational mechanism by which choice-induced preference changes might occur. In their comparison of evidence-integration models of choice, Glickman et al. (2022) showed that the most successful model was one that overweighted new evidence when it was consistent with immediately preceding evidence. Surprisingly, such apparent distortion of evidence led to choices with higher accuracy as well as higher confidence. Although the distortion discussed in their study was related to the difference in value between options, it is possible that a similar distortion might take place within the individual value signals (for each option) before they are compared. Such distortions might provide a benefit to the decision system by dampening the impact of neural noise, causing information about the positive and negative attributes of each option to be accentuated. Future work will be needed to directly test this intriguing hypothesis.

Future work might also investigate whether the rate at which SoA dissipates after choice may vary as a function of option attribute disparity (i.e., the magnitude of differences in the attribute compositions of the choice options). Previous work has shown that SoA is positively correlated with disparity (Lee & Holyoak, 2021). It could be that SoA is driven by different processes when disparity is high versus low (e.g., revisions of attribute importance weights vs. revisions in attribute measurements). It would be interesting to examine whether the rate of SoA dissipation varies according to the attribute compositions of the options.

References

Alós-Ferrer, C., & Shi, F. (2015). Choice-induced preference change and the free-choice paradigm: A clarification. *Judgment and Decision Making*, *10*(1), 34–49. https://doi.org/10.1017/S1930297500003168

Bem, D. J. (1967). Self-perception: An alternative interpretation of cognitive dissonance phenomena. *Psychological Review*, *74*(3), 183–200. https://doi.org/10.1037/h0024835

Bem, D. J. (1972). Self-perception theory. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 6, pp. 1–62). Academic Press. https://doi.org/10.1016/S0065-2601(08)60024-6

Brehm, J. W. (1956). Postdecision changes in the desirability of alternatives. *Journal of Abnormal Psychology*, *52*(3), 384–389. https://doi.org/10.1037/h0041006

Brownstein, A. L. (2003). Biased predecision processing. *Psychological Bulletin*, *129*(4), 545–568. https://doi.org/10.1037/0033-2909.129.4.545

Carlson, K. A., Meloy, M. G., & Russo, J. E. (2006). Leader-driven primacy: Using attribute order to affect consumer choice. *The Journal of Consumer Research*, *32*(4), 513–518. https://doi.org/10.1086/500481

Carlson, K. A., & Russo, J. E. (2001). Biased interpretation of evidence by mock jurors. *Journal of Experimental Psychology: Applied*, *7*(2), 91–103. https://doi.org/10.1037/1076-898X.7.2.91

Chammat, M., Karoui, I. E., 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), Article 41320. https://doi.org/10.1038/srep41320

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. https://doi.org/10.1037/a0020217

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. https://doi.org/10.1177/0963721415587876

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. https://doi.org/10.3758/s13423-011-0184-8

Engel, C., & Glöckner, A. (2013). Role-induced bias in court: An experimental analysis. *Journal of Behavioral Decision Making*, *26*(3), 272–284. https://doi.org/10.1002/bdm.1761

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. https://doi.org/10.1037/pspa0000263

Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford University Press. https://doi.org/10.1515/9781503620766

Glickman, M., Moran, R., & Usher, M. (2022). Evidence integration and decision confidence are modulated by stimulus consistency. *Nature Human Behaviour*, *6*(7), 988–999. https://doi.org/10.1038/s41562-022-01318-6

Glöckner, A., Betsch, T., & Schindler, N. (2010). Coherence shifts in probabilistic inference tasks. *Journal of Behavioral Decision Making*, *23*(5), 439–462. https://doi.org/10.1002/bdm.668

Harmon-Jones, E., & Mills, J. (2019). An introduction to cognitive dissonance theory and an overview of current perspectives on the theory. In E. Harmon-Jones (Ed.), *Cognitive dissonance: Reexamining a pivotal theory in psychology* (2nd ed., pp. 3–24). American Psychological Association. https://doi.org/10.1037/0000135-001

Holyoak, K. J., & Simon, D. (1999). Bidirectional reasoning in decision making by constraint satisfaction. *Journal of Experimental Psychology: General*, *128*(1), 3–31. https://doi.org/10.1037/0096-3445.128.1.3

Izuma, K., Akula, S., Murayama, K., Wu, D.-A., Iacoboni, M., & Adolphs, R. (2015). A causal role for posterior medial frontal cortex in choice-induced preference change. *The Journal of Neuroscience*, *35*(8), 3598–3606. https://doi.org/10.1523/JNEUROSCI.4591-14.2015

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. https://doi.org/10.1073/pnas.1011879108

Izuma, K., & Murayama, K. (2013). Choice-induced preference change in the free-choice paradigm: A critical methodological review. *Frontiers in Psychology*, *4*, Article 41. https://doi.org/10.3389/fpsyg.2013.00041

Kostopoulou, O., Russo, J. E., Keenan, G., Delaney, B. C., & Douiri, A. (2012). Information distortion in physicians' diagnostic judgments. *Medical Decision Making*, *32*(6), 831–839. https://doi.org/10.1177/0272989X12447241

Lee, D. G. (2021). *Lee & Holyoak, 2023—Transient value refinements during deliberation facilitate choice*. https://doi.org/10.17605/OSF.IO/PQW4X

Lee, D. G., & Coricelli, G. (2020). An empirical test of the role of value certainty in decision making. *Frontiers in Psychology*, *11*, Article 574473. https://doi.org/10.3389/fpsyg.2020.574473

Lee, D. G., & 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. https://doi.org/10.1371/journal.pone.0231081

Lee, D. G., & Daunizeau, J. (2021). Trading mental effort for confidence in the metacognitive control of value-based decision-making. *eLife*, *10*, Article e63282. https://doi.org/10.7554/eLife.63282

Lee, D. G., & Hare, T. A. (2023). Value certainty and choice confidence are multidimensional constructs that guide decision-making. *Cognitive, Affective & Behavioral Neuroscience*. Advance online publication. https://doi.org/10.3758/s13415-022-01054-4

Lee, D. G., & Holyoak, K. J. (2021). Coherence shifts in attribute evaluations. *Decision*, *8*(4), 257–276. https://doi.org/10.1037/dec0000151

Lee, D. G., & Pezzulo, G. (2022). *Changes in preferences reported after choices are informative, not merely statistical artifacts*. PsyArXiv. https://doi.org/10.31234/osf.io/btj95

Miller, S. A., DeKay, M. L., Stone, E. R., & Sorenson, C. M. (2013). Assessing the sensitivity of information distortion to four potential influences in studies of risky choice. *Judgment and Decision Making*, *8*(6), 662–677. https://doi.org/10.1017/S1930297500004708

Nurek, M., Kostopoulou, O., & Hagmayer, Y. (2014). Predecisional information distortion in physicians' diagnostic judgments: Strengthening a leading hypothesis or weakening its competitor? *Judgment and Decision Making*, *9*(6), 572–585. https://doi.org/10.1017/S1930297500006434

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. https://doi.org/10.1037/a0012786

Russo, J. E., Medvec, V. H., & Meloy, M. G. (1996). The distortion of information during decisions. *Organizational Behavior and Human Decision Processes*, *66*(1), 102–110. https://doi.org/10.1006/obhd.1996.0041

Russo, J. E., Meloy, M. G., & Medvec, V. H. (1998). Predecisional distortion of product information. *Journal of Marketing Research*, *35*(4), 438–452. https://doi.org/10.1177/002224379803500403

Salti, M., El Karoui, I., Maillet, M., & Naccache, L. (2014). Cognitive dissonance resolution is related to episodic memory. *PLOS ONE*, *9*(9), Article e108579. https://doi.org/10.1371/journal.pone.0108579

Schonberg, T., & Katz, L. N. (2020). A neural pathway for nonreinforced preference change. *Trends in Cognitive Sciences*, *24*(7), 504–514. https://doi.org/10.1016/j.tics.2020.04.002

Sharot, T., De Martino, B., & Dolan, R. J. (2009). How choice reveals and shapes expected hedonic outcome. *The Journal of Neuroscience*, *29*(12), 3760–3765. https://doi.org/10.1523/JNEUROSCI.4972-08.2009

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. https://doi.org/10.1177/0956797612438733

Sharot, T., Velasquez, C. M., & Dolan, R. J. (2010). Do decisions shape preference? Evidence from blind choice. *Psychological Science*, *21*(9), 1231–1235. https://doi.org/10.1177/0956797610379235

Simon, D., Krawczyk, D. C., Bleicher, A., & Holyoak, K. J. (2008). The transience of constructed preferences. *Journal of Behavioral Decision Making*, *21*(1), 1–14. https://doi.org/10.1002/bdm.575

Simon, D., Krawczyk, D. C., & Holyoak, K. J. (2004). Construction of preferences by constraint satisfaction. *Psychological Science*, *15*(5), 331–336. https://doi.org/10.1111/j.0956-7976.2004.00678.x

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*(5), 1250–1260. https://doi.org/10.1037/0278-7393.27.5.1250

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), 814–837. https://doi.org/10.1037/0022-3514.86.6.814

Simon, D., & Spiller, S. A. (2016). The elasticity of preferences. *Psychological Science*, *27*(12), 1588–1599. https://doi.org/10.1177/0956797616666501

Talluri, B. C., Urai, A. E., Tsetsos, K., Usher, M., & Donner, T. H. (2018). Confirmation bias through selective overweighting of choice-consistent evidence. *Current Biology*, *28*(19), 3128–3135.e8. https://doi.org/10.1016/j.cub.2018.07.052

Voigt, K., Murawski, C., & Bode, S. (2017). Endogenous formation of preferences: Choices systematically change willingness-to-pay for goods. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *43*(12), 1872–1882. https://doi.org/10.1037/xlm0000415

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. https://doi.org/10.1523/JNEUROSCI.1681-18.2018

Received January 5, 2022
Revision received March 8, 2023
Accepted March 10, 2023 ■