Transient Value Refinements during Deliberation Facilitate Choice

Douglas G. Lee ^a

Keith J. Holyoak ^b

^a Tel Aviv University

School of Psychological Sciences

Tel Aviv 6997801

Israel

douglas@tauex.tau.ac.il

^b University of California, Los Angeles

Department of Psychology

405 Hilgard Avenue

Los Angeles, CA 90095-1563

USA

kholyoak@g.ucla.edu

Contact info for correspondence:

DouglasGLee@gmail.com

Tel Aviv University

School of Psychological Sciences

Tel Aviv 6997801

Israel

Abstract

After deciding among options, decision-makers tend to exhibit an increased evaluation of the chosen options coupled with a decreased evaluation of the rejected options, resulting in a *spreading of alternatives* (SoA). There has been a longstanding debate as to whether SoA results from post-choice 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 enhances choice consistency, boosts confidence, and reduces response time. The SoA generated during choice deliberation is partially transient in nature, regressing toward baseline shortly afterwards, but lingering at an intermediate level.

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

Statement of Relevance

Contrary to the conventional notion that choices are based on pre-existing values of options, people may sometimes construct these values in the very process of making a choice. In addition to values determining choices, choices may partially determine values, resulting in a *spreading of alternatives*. A longstanding debate has centered on whether the spreading of alternatives solely reflects post hoc justification of a decision already made, or whether it impacts how a decision emerges. Using a novel method to measure the spreading of alternatives during an implicit choice task, we show that option values change *before* a choice is reported. Moreover, this process of value refinement is instrumental to making a choice: it enhances choice consistency, boosts confidence, and reduces response time. Once the choice has been made, option values partially regress toward baseline. These findings help to understand how people make consequential decisions, ranging from choice of consumer products to deciding legal cases.

1. Introduction

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: pre-choice ratings of individual options; choices between pairs of options; and post-choice 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 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 afterwards 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 post-choice 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 (Lee & Daunizeau, 2021) proposed a model that explains SoA based on three assumptions: 1) pre-choice value estimates are imprecise; 2) choices between similarly-valued options encourage refinement of value estimates; 3) post-choice 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 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 pre-existing 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, post-choice ratings of choice option pairs will be (on average) farther apart relative to the corresponding pre-choice ratings. It has been demonstrated via simulation that the statistical observation holds when the difference in pre-choice ratings is sufficiently small ((Alós-Ferrer & Shi, 2015; Izuma & Murayama, 2013; Lee & Pezzulo, 2022); see also Supplementary 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 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 pre-decisional 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 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 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 minutes 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 increased choice consistency and confidence, and reduced response time; SoA then dissipated after the choice, with value estimates regressing (though only partially) back towards 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.

2. Material and methods

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 post-choice ratings.

2.1 *Participants*

A total of 178 people participated (90 female; age: mean = 42 years, SD = 8, range = 24-55), split evenly across three experimental conditions (*Standard Choice*: $n = 60$, 28 female, age: mean = 41 years, SD = 8, range = 26-55; *Implicit Choice*: $n = 60$, 30 female, age: mean = 41 years, SD = 8, range = 24-55; *Standard + Implicit Choice*: $n = 58$, 32 female, age: mean = 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 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 minutes 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.

2.2 *Materials*

The experiments were constructed using the online experimental platform Gorilla (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.

2.3 *Design and Procedure*

The experiment consisted of four phases: pre-exposure, 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

168 of snack foods. Specifically, for each of a variety of snacks, you will tell us how much you feel
169 that you would enjoy eating it on a frequent basis. You will also be asked to choose your preferred
170 item from different pairs, as well as to estimate how confident you are about each choice.” The
171 instructions that were presented to participants at the beginning of each task are provided in
172 Supplementary Material.

Figure 1: Experimental design. In the Exposure task, participants passively viewed each option for 1 second. 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.

In the pre-exposure phase, participants simply observed as all individual items were displayed in a random sequence for 1000 ms each (with no inter-trial interval). 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 value across the item set. Just prior to the onset of the pre-exposure 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, 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 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 current 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 three conditions for this task, counterbalanced across conditions. In all three conditions, 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 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.”

2.4 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 outlier 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.

2.5 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 (non-regression) results, reported p values represent the probability of non-zero effect sizes, 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 post-choice), we distinguish the rating difference from these phases by labeling them $dV1$ (pre-choice) and $dV2$ (post-choice). The Implicit Choice and Standard + Implicit Choice conditions also involved an additional rating phase (intra-choice). 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 spreading of alternatives 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 $\text{SoA} = \text{SoA}^* + \text{SoA}^R$.

Apart from SoA, we also examine choice consistency (defined as a choice in favor of the option that was rated with a higher value), response time (RT; measured from the presentation of

the choice options until 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 co-regressor 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 co-regressor 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).

2.6 Ruling Out the Statistical Artifact Explanation

Chen and Risen (Chen & Risen, 2010) pointed out that a statistical artifact in the way in which the spreading of alternatives (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 post-choice. 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 (Lee & 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 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 Supplementary Material.

3. 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 seconds, 5.4 seconds, and 5.4 seconds 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, 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 pre-choice 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 non-directional, 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 pre-choice ratings but consistent with post-choice rating) should be more common than other types of inconsistency (i.e., choice consistent with pre- but not post-choice ratings – “rating error” – or consistent ratings with 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; Table 1)

Table 1. Comparisons of changes of mind versus rating and choice errors.

Standard Condition					
DV	M	C.I.	t	d.f.	p
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	0.011
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					
DV	M	C.I.	t	d.f.	p
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					
DV	M	C.I.	t	d.f.	p
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	0.011

CoM > err _r (difficult only)	0.052	[0.0135 0.0905]	2.703	59	0.009
CoM > err _c (difficult only)	0.080	[0.0400 0.1205]	3.988	59	< .001

3.1 *Choice Induces Preference Change*

In accord with previous work, we observed a reliable SoA³ across all choice trials (cross-participant mean of 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 $p > 0.727$), 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).

Table 2. Regression coefficients relating choice ease and SoA, across experimental conditions.

Standard					
DV	M	S.E.	t	d.f.	p
$\beta dV1 $ on SoA	-0.13	0.02	-7.51	2859	< .001
Standard + Implicit					
DV	M	S.E.	t	d.f.	p
$\beta dV1 $ on SoA	-0.21	0.02	-10.19	2767	< .001
Implicit					
DV	M	S.E.	t	d.f.	p
$\beta dV1 $ on SoA	-0.21	0.02	-10.72	2886	< .001

3.2 *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 post-choice ratings (i.e., dV2) rather than pre-choice 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 current study replicates this finding. Specifically, we first regressed consistency on dV1 and dV2 residuals (see Methods). For the Standard + Implicit and Implicit conditions, we also included dV* residuals (see Methods) as a regressor. If intra-choice 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 post-choice 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 Table 3). 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 Table 3).

Figure 2. 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.

Table 3. Regression coefficients relating value difference and choice consistency, across experimental conditions.

Standard					
DV	M	S.E.	t	d.f.	p
β dV1 on Consistency	1.06	0.03	30.55	2858	< .001
β dV2 _{resid} on Consistency	0.79	0.06	13.01	2858	< .001
Standard + Implicit					
DV	M	S.E.	t	d.f.	p
β dV1 on Consistency	1.15	0.05	20.95	2765	< .001
β dV2 _{resid} on Consistency	0.81	0.05	17.06	2765	< .001
β dV* _{resid} on Consistency	0.87	0.06	14.77	2765	< .001
Implicit					
DV	M	S.E.	t	d.f.	p
β dV1 on Consistency	1.11	0.05	23.97	2884	< .001
β dV2 _{resid} on Consistency	0.93	0.05	17.97	2884	< .001
β dV* _{resid} on Consistency	1.35	0.07	18.05	2884	< .001

Previous studies have shown that SoA positively influences 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 pre-choice 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 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, since it represents changes in value estimates that occur

407 after the choices are reported. Or, SoA^R might have a negative predictive effect, if it was merely
408 caused by noise unrelated to choice. Finally, SoA^R might have a positive predictive effect, if it
409 reflects a sort of correction against exaggerated coherence shifts (i.e., SoA^* resulting in dV^*
410 greater than the true subjective value difference). Unsurprisingly, we found that $|dV1|$ has a reliable
411 positive relationship with confidence and a reliable negative relationship with RT (Figure 3 and
412 Tables 4-5). Critically, for both dependent variables in all conditions, we found that SoA^* has an
413 incremental predictive effect beyond the effect of $|dV1|$ (Figure 3 and Tables 4-5). Interestingly,
414 for both dependent variables in both relevant conditions, we also found that SoA^R has an
415 incremental predictive effect beyond the effect of $|dV1|$ and SoA^* (Figure 3 and Tables 4-5). All
416 beta weights for SoA resembled those for $|dV1|$ (though with smaller magnitudes), demonstrating
417 that those variables serve a similar role in determining choice behavior.

Figure 3. Predictive effect of |dV1|, SoA*, and SoA^R on RT (left plots) and confidence (right plots), across conditions (rows; recall that the separation of SoA into SoA* + SoA^R does not apply in the Standard condition). Error bars represent 95% confidence intervals.

Table 4. Regression coefficients relating choice ease and SoA with RT, across experimental conditions.

Standard					
DV	M	S.E.	t	d.f.	p
β dV1 on RT	-0.59	0.04	-14.40	2858	< .001
β SoA on RT	-0.20	0.05	-4.03	2858	< .001
Standard + Implicit					
DV	M	S.E.	t	d.f.	p
β dV1 on RT	-0.25	0.04	-6.73	2765	< .001
β SoA* on RT	-0.18	0.04	-4.41	2765	< .001
β SoA ^R on RT	-0.06	0.04	-1.35	2765	0.176

Implicit					
DV	M	S.E.	t	d.f.	p
β dV1 on RT	-0.13	0.03	-3.91	2884	< .001
β SoA* on RT	-0.08	0.05	-1.85	2884	0.065
β SoA ^R on RT	-0.09	0.04	-2.21	2884	0.027

Table 5. Regression coefficients relating choice ease and SoA with confidence, across experimental conditions.

Standard					
DV	M	S.E.	t	d.f.	p
β dV1 on Confidence	0.54	0.04	14.85	2858	< .001
β SoA on Confidence	0.17	0.03	5.95	2858	< .001
Standard + Implicit					
DV	M	S.E.	t	d.f.	p
β dV1 on Confidence	0.49	0.05	10.76	2765	< .001
β SoA* on Confidence	0.35	0.04	7.99	2765	< .001
β SoA ^R on Confidence	0.13	0.03	4.39	2765	< .001
Implicit					
DV	M	S.E.	t	d.f.	p
β dV1 on Confidence	0.35	0.05	7.43	2884	< .001
β SoA* on Confidence	0.19	0.04	4.94	2884	< .001
β SoA ^R on Confidence	0.12	0.03	4.11	2884	< .001

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 more predictive power than dV1, as the post-choice ratings should include more decision-relevant information than the pre-choice ratings. If post-choice ratings are richer in decision-relevant information compared to pre-choice 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 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, adjusted r -squared was larger for the regressions based on $dV2$ (Table 6).

Table 6: Post-choice ratings explain choice data better than pre-choice ratings. For all dependent variables across all conditions, a larger percentage of variance in the data was explained by models based on post-choice ratings ($dV2$) versus pre-choice ratings ($dV1$).

	R^2	
	$dV1$	$dV2$
<u>Standard</u>		
consistency	0.43	0.48
RT	0.36	0.36
confidence	0.36	0.39
<u>Implicit</u>		
consistency	0.47	0.52
RT	0.36	0.36
confidence	0.36	0.41
<u>Standard + Implicit</u>		
consistency	0.44	0.50
RT	0.43	0.44
confidence	0.45	0.49

3.3 *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: spreading of alternatives *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 post-choice 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 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 towards 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).

Figure 4. Value difference estimates are exaggerated during choice deliberation. An illustration of how choice-induced preference change might arise. A) Pre-choice ratings show an estimated value difference (dV1) of 10. B) Intra-choice ratings that inform the decision itself show an exaggerated value difference (dV*) of 30. C) Post-choice ratings show an estimated value difference (dV2) of 20. In this extreme (for illustrative purposes) example, SoA* is 20 and final SoA is 10, with a transiency of 50% as in the experimental data.

To quantify the transient nature, we separated 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 (Figure 5).

Figure 5. Transience of Spreading of Alternatives (SoA). 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 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; red lines represent cross-participant median values.

4. Discussion

The spreading of alternatives (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 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 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 intra-decision 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 pre-decisional 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. 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 post-choice rather than intra-choice. 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 strong 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 post-choice 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 pre-choice 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 enhanced choice consistency. It appears that the SoA observed in the present study truly made choices easier (and thus was not random), as SoA decreased response time and increased choice confidence (Lee & Daunizeau, 2020, 2021). Moreover, post-choice ratings had higher predictive power on choice consistency, RT, and confidence than did pre-choice ratings, which would not be the case if the rating changes were a mere statistical artifact. 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 Supplementary Material).

It might be argued that memory may have played some role in causing post-choice ratings to be more consistent with choices than pre-choice ratings. One might claim that when rating an option post-choice, a decision-maker simply adjusts the pre-choice 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 minutes prior, and 50 previous choices made up to 20 minutes prior (in totally random sequences). A similar alternative explanation might suggest that post-choice 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 post-choice ratings imply preferences opposite to those implied by pre-choice 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 unequivocally rule out the possibility that participants might first decide (in their minds) which option to choose, 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 developing 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 intra-choice ratings) was substantially greater than final SoA (based on the change in dV from pre- to post-choice 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 post-choice ratings) as did previous studies, the present study is the first to separate SoA into pre- and post-choice SoA. We found that post-choice SoA (between choice and post-choice ratings) was often *negative*, and that the magnitude of negative post-choice SoA was strongly correlated with the magnitude of pre-choice SoA. This pattern suggests that while 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 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 and colleagues (Glickman et al., 2022) showed that the best model was one that over-weighted 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. While 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 (Glickman et al., 2022), 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 versus revisions in attribute measurements). It would be interesting to examine whether rate of SoA dissipation varies according to the attribute compositions of the options.

646

Footnote

647

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).

650

651

652

2. For consistent choices, $\text{SoA} = dV2 - dV1$; for changes of mind, $\text{SoA} = dV1 - dV2$.

653

Author Contributions

654 D.L. developed the study concept. D.L. and K.H. contributed to the study design. D.L. performed
655 data collection and analysis. D.L. drafted the manuscript, D.L. and K.H. provided critical revisions.
656 Both authors approved the final version of the manuscript for submission.

657

Funding

658 This work was supported by the National Science Foundation [grant number BCS-1827374].

659

660

Data and Analysis Code Availability

661 The raw data and primary analysis code for this study are available in the Open Science Forum at
662 <http://doi.org/10.17605/OSF.IO/PQW4X>.

References

- Alós-Ferrer, C., & Shi, F. (2015). Choice-induced preference change and the free-choice paradigm: A clarification. *Judgment and Decision Making*, 10(1), 16.
- 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](https://doi.org/10.1016/S0065-2601(08)60024-6)
- Brehm, J. W. (1956). Postdecision changes in the desirability of alternatives. *The Journal of Abnormal and Social 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. *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. <https://doi.org/10.1038/srep41320>

- 685 Chen, K. M., & Risen, J. L. (2010). How choice affects and reflects preferences: Revisiting the
686 free-choice paradigm. *Journal of Personality and Social Psychology*, 99(4), 573–594.
687 <https://doi.org/10.1037/a0020217>
- 688 DeKay, M. L. (2015). Predecisional Information Distortion and the Self-Fulfilling Prophecy of
689 Early Preferences in Choice. *Current Directions in Psychological Science*, 24(5), 405–
690 411. <https://doi.org/10.1177/0963721415587876>
- 691 DeKay, M. L., Stone, E. R., & Sorenson, C. M. (2012). Sizing up information distortion:
692 Quantifying its effect on the subjective values of choice options. *Psychonomic Bulletin &*
693 *Review*, 19(2), 349–356. <https://doi.org/10.3758/s13423-011-0184-8>
- 694 Engel, C., & Glöckner, A. (2013). Role-Induced Bias in Court: An Experimental Analysis.
695 *Journal of Behavioral Decision Making*, 26(3), 272–284.
696 <https://doi.org/10.1002/bdm.1761>
- 697 Enisman, M., Shpitzer, H., & Kleiman, T. (2021). Choice changes preferences, not merely
698 reflects them: A meta-analysis of the artifact-free free-choice paradigm. *Journal of*
699 *Personality and Social Psychology*, 120(1), 16–29. <https://doi.org/10.1037/pspa0000263>
- 700 Festinger, L. (1957). *A Theory of Cognitive Dissonance*. Stanford University Press.
- 701 Glickman, M., Moran, R., & Usher, M. (2022). Evidence integration and decision confidence are
702 modulated by stimulus consistency. *Nature Human Behaviour*, 6(7), Article 7.
703 <https://doi.org/10.1038/s41562-022-01318-6>
- 704 Glöckner, A., Betsch, T., & Schindler, N. (2010). Coherence shifts in probabilistic inference
705 tasks. *Journal of Behavioral Decision Making*, 23(5), 439–462.
706 <https://doi.org/10.1002/bdm.668>

- 707 Harmon-Jones, E., & Mills, J. (2019). An introduction to cognitive dissonance theory and an
708 overview of current perspectives on the theory. In *Cognitive dissonance: Reexamining a*
709 *pivotal theory in psychology, 2nd ed* (pp. 3–24). American Psychological Association.
710 <https://doi.org/10.1037/0000135-001>
- 711 Holyoak, K. J., & Simon, D. (1999). Bidirectional reasoning in decision making by constraint
712 satisfaction. *Journal of Experimental Psychology: General*, 128(1), 3–31.
713 <https://doi.org/10.1037/0096-3445.128.1.3>
- 714 Izuma, K., Akula, S., Murayama, K., Wu, D.-A., Iacoboni, M., & Adolphs, R. (2015). A causal
715 role for posterior medial frontal cortex in choice-induced preference change. *The Journal*
716 *of Neuroscience: The Official Journal of the Society for Neuroscience*, 35(8), 3598–3606.
717 <https://doi.org/10.1523/JNEUROSCI.4591-14.2015>
- 718 Izuma, K., Matsumoto, M., Murayama, K., Samejima, K., Sadato, N., & Matsumoto, K. (2010).
719 Neural correlates of cognitive dissonance and choice-induced preference change.
720 *Proceedings of the National Academy of Sciences*, 107(51), 22014–22019.
721 <https://doi.org/10.1073/pnas.1011879108>
- 722 Izuma, K., & Murayama, K. (2013). Choice-Induced Preference Change in the Free-Choice
723 Paradigm: A Critical Methodological Review. *Frontiers in Psychology*, 4.
724 <https://doi.org/10.3389/fpsyg.2013.00041>
- 725 Kostopoulou, O., Russo, J. E., Keenan, G., Delaney, B. C., & Douiri, A. (2012). Information
726 Distortion in Physicians' Diagnostic Judgments. *Medical Decision Making*, 32(6), 831–
727 839. <https://doi.org/10.1177/0272989X12447241>
- 728 Lee, D. G., & Coricelli, G. (2020). An Empirical Test of the Role of Value Certainty in Decision
729 Making. *Frontiers in Psychology*, 11, 574473. <https://doi.org/10.3389/fpsyg.2020.574473>

- 730 Lee, D. G., & Daunizeau, J. (2020). Choosing what we like vs liking what we choose: How
731 choice-induced preference change might actually be instrumental to decision-making.
732 *PLOS ONE*, 15(5), e0231081. <https://doi.org/10.1371/journal.pone.0231081>
- 733 Lee, D. G., & Daunizeau, J. (2021). Trading mental effort for confidence in the metacognitive
734 control of value-based decision-making. *ELife*, 10, e63282.
735 <https://doi.org/10.7554/eLife.63282>
- 736 Lee, D. G., & Hare, T. A. (2023). Value certainty and choice confidence are multidimensional
737 constructs that guide decision-making. *Cognitive, Affective, & Behavioral Neuroscience*.
738 <https://doi.org/10.3758/s13415-022-01054-4>
- 739 Lee, D. G., & Holyoak, K. J. (2021). Coherence shifts in attribute evaluations. *Decision*, 8(4),
740 257–276. <https://doi.org/10.1037/dec0000151>
- 741 Lee, D. G., & Pezzulo, G. (2022). *Changes in Preferences Reported After Choices Are*
742 *Informative, Not Merely Statistical Artifacts*. PsyArXiv.
743 <https://doi.org/10.31234/osf.io/btj95>
- 744 Miller, S. A., DeKay, M. L., Stone, E. R., & Sorenson, C. M. (2013). Assessing the sensitivity of
745 information distortion to four potential influences in studies of risky choice. *Judgment*
746 *and Decision Making*, 8(6), 662–677. <https://doi.org/10.1017/S1930297500004708>
- 747 Nurek, M., Kostopoulou, O., & Hagmayer, Y. (2014). Predecisional information distortion in
748 physicians' diagnostic judgments: Strengthening a leading hypothesis or weakening its
749 competitor? *Judgment and Decision Making*, 9(6).
750 [https://kclpure.kcl.ac.uk/portal/en/publications/predecisional-information-distortion-in-](https://kclpure.kcl.ac.uk/portal/en/publications/predecisional-information-distortion-in-physicians-diagnostic-judgments(280e11f4-044a-49d5-8dd2-f5572a37b0d9).html)
751 [physicians-diagnostic-judgments\(280e11f4-044a-49d5-8dd2-f5572a37b0d9\).html](https://kclpure.kcl.ac.uk/portal/en/publications/predecisional-information-distortion-in-physicians-diagnostic-judgments(280e11f4-044a-49d5-8dd2-f5572a37b0d9).html)

- 752 Pleskac, T. J., & Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of
753 choice, decision time, and confidence. *Psychological Review*, 117(3), 864–901.
754 <https://doi.org/10.1037/a0019737>
- 755 Russo, J. E. (2015). The predecisional distortion of information. In *Neuroeconomics, judgment,*
756 *and decision making* (pp. 91–110). Psychology Press.
- 757 Russo, J. E., Carlson, K. A., Meloy, M. G., & Yong, K. (2008). The goal of consistency as a
758 cause of information distortion. *Journal of Experimental Psychology: General*, 137(3),
759 456–470. <https://doi.org/10.1037/a0012786>
- 760 Russo, J. E., Medvec, V. H., & Meloy, M. G. (1996). The Distortion of Information during
761 Decisions. *Organizational Behavior and Human Decision Processes*, 66(1), 102–110.
762 <https://doi.org/10.1006/obhd.1996.0041>
- 763 Russo, J. E., Meloy, M. G., & Medvec, V. H. (1998). Predecisional Distortion of Product
764 Information. *Journal of Marketing Research*, 35(4), 438–452.
765 <https://doi.org/10.1177/002224379803500403>
- 766 Salti, M., El Karoui, I., Maillet, M., & Naccache, L. (2014). Cognitive Dissonance Resolution Is
767 Related to Episodic Memory. *PLoS ONE*, 9(9).
768 <https://doi.org/10.1371/journal.pone.0108579>
- 769 Schonberg, T., & Katz, L. N. (2020). A Neural Pathway for Nonreinforced Preference Change.
770 *Trends in Cognitive Sciences*, 24(7), 504–514. <https://doi.org/10.1016/j.tics.2020.04.002>
- 771 Sharot, T., De Martino, B., & Dolan, R. J. (2009). How choice reveals and shapes expected
772 hedonic outcome. *The Journal of Neuroscience: The Official Journal of the Society for*
773 *Neuroscience*, 29(12), 3760–3765. <https://doi.org/10.1523/JNEUROSCI.4972-08.2009>

- 774 Sharot, T., Fleming, S. M., Yu, X., Koster, R., & Dolan, R. J. (2012). Is Choice-Induced
775 Preference Change Long Lasting? *Psychological Science*, 23(10), 1123–1129.
776 <https://doi.org/10.1177/0956797612438733>
- 777 Sharot, T., Velasquez, C. M., & Dolan, R. J. (2010). Do Decisions Shape Preference? Evidence
778 From Blind Choice. *Psychological Science*, 21(9), 1231–1235.
- 779 Simon, D., Krawczyk, D. C., Bleicher, A., & Holyoak, K. J. (2008). The transience of
780 constructed preferences. *Journal of Behavioral Decision Making*, 21(1), 1–14.
781 <https://doi.org/10.1002/bdm.575>
- 782 Simon, D., Krawczyk, D. C., & Holyoak, K. J. (2004). Construction of preferences by constraint
783 satisfaction. *Psychological Science*, 15(5), 331–336. [https://doi.org/10.1111/j.0956-](https://doi.org/10.1111/j.0956-7976.2004.00678.x)
784 [7976.2004.00678.x](https://doi.org/10.1111/j.0956-7976.2004.00678.x)
- 785 Simon, D., Pham, L. B., Le, Q. A., & Holyoak, K. J. (2001). The emergence of coherence over
786 the course of decision making. *Journal of Experimental Psychology: Learning, Memory,*
787 *and Cognition*, 27(5), 1250–1260. <https://doi.org/10.1037/0278-7393.27.5.1250>
- 788 Simon, D., Snow, C. J., & Read, S. J. (2004). The Redux of Cognitive Consistency Theories:
789 Evidence Judgments by Constraint Satisfaction. *Journal of Personality and Social*
790 *Psychology*, 86(6), 814–837. <https://doi.org/10.1037/0022-3514.86.6.814>
- 791 Simon, D., & Spiller, S. A. (2016). The Elasticity of Preferences. *Psychological Science*, 27(12),
792 1588–1599. <https://doi.org/10.1177/0956797616666501>
- 793 Talluri, B. C., Urai, A. E., Tsetsos, K., Usher, M., & Donner, T. H. (2018). Confirmation Bias
794 through Selective Overweighting of Choice-Consistent Evidence. *Current Biology*,
795 28(19), 3128–3135.e8. <https://doi.org/10.1016/j.cub.2018.07.052>

- 796 Vickers, D., & Packer, J. (1982). Effects of alternating set for speed or accuracy on response
797 time, accuracy and confidence in a unidimensional discrimination task. *Acta*
798 *Psychologica*, 50(2), 179–197. [https://doi.org/10.1016/0001-6918\(82\)90006-3](https://doi.org/10.1016/0001-6918(82)90006-3)
- 799 Voigt, K., Murawski, C., & Bode, S. (2017). Endogenous formation of preferences: Choices
800 systematically change willingness-to-pay for goods. *Journal of Experimental Psychology*
801 *Learning, Memory, and Cognition*, 43(12), 1872–1882.
802 <https://doi.org/10.1037/xlm0000415>
- 803 Voigt, K., Murawski, C., Speer, S., & Bode, S. (2019). Hard Decisions Shape the Neural Coding
804 of Preferences. *Journal of Neuroscience*, 39(4), 718–726.
805 <https://doi.org/10.1523/JNEUROSCI.1681-18.2018>
806

Supplementary Material

Rejecting the Statistical Artifact Explanation

Chen and Risen (Chen & Risen, 2010) pointed out that a statistical artifact in the way in which the spreading of alternatives (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 post-choice. Lee and Pezzulo (Lee & 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. 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.

To undertake the necessary analyses, we generated a set of synthetic data using the actual empirical measures of the value estimates for each choice option (averaged across pre-and post-choice ratings). We first imported the value estimates provided by all participants across all conditions ($n_{\text{Trials}} = 8900$), calculated as the average of pre- and post-choice ratings for each option. We repeated each trial 10 times to reduce the effect of simulation noise ($n = 89000$). We then created a set of pre-choice ratings for all options by sampling from normal distributions where the means were equal to the empirical value estimates and the standard deviation was set at 0.1 (different levels of noise are discussed in (Lee & Pezzulo, 2022)). Subtracting the value of option 2 from the value of option 1 on each trial provided a measure of pre-choice value difference ($dV1$).

For the *statistical artifact* (Chen & Risen, 2010) simulation, we determined choices deterministically, where the option with a higher “true” value (from the empirical input) within

each pair was considered the chosen option on that trial. By eliminating noise from the choice process, we maximize the opportunity for SoA to arise (Izuma & Murayama, 2013) and thus give this account the greatest possibility of explaining the data. We then created a set of post-choice ratings for all options in exactly the same way as for the pre-choice ratings, which provided a measure of post-choice value difference (dV_2). We calculated SoA in the same manner as with our empirical data: the post- minus pre-choice ratings of the chosen options minus the post- minus pre-choice ratings of the rejected options. The statistical artifact account of SoA does not inherently include RT or choice confidence. Nevertheless, we followed Lee and Pezzulo (Lee & Pezzulo, 2022) and used a heuristic approach to calculate confidence as a sigmoidal function of the difference in the pre-choice values of the options on each trial (though any monotonically-increasing function would have been acceptable). We then calculated RT as the inverse of confidence (since confidence and RT are known to be inversely related, empirically; (Pleskac & Busmeyer, 2010; Vickers & Packer, 1982)).

We also simulated a set of data under an account where SoA is caused by *value refinement* and *precision gain* thought to occur during choice deliberation (see (Lee & Daunizeau, 2021)). We used the same pre-choice ratings under this account as we used under the statistical artifact account. For choices, we first perturbed the true value of each option by a random amount inversely proportional to the pre-choice difficulty: $perturbation\ factor * (1 - \text{abs}(dV_1)) * \text{uniform}(-1,1)$. We set the perturbation factor to 0.1 (different levels of perturbation factor are discussed in (Lee & Pezzulo, 2022)). Choices were again deterministic, based on the refined true values. We then created a set of post-choice ratings for all options in a similar manner as for the pre-choice ratings but based on the refined true values and including a gain in precision (relative to the pre-choice ratings) inversely proportional to the pre-choice difficulty: $precision\ factor * (1 - \text{abs}(dV_1))$.

Specifically, the standard deviations of the normal distributions (with means equal to the perturbed true values) were divided by $(1 + \text{gain})$, with the gain being specific to each trial. We set the precision factor to 0.3 (different levels of precision factor are discussed in (Lee & Pezzulo, 2022)). Our simulations were consistent with the Lee & Daunizeau (Lee & Daunizeau, 2021) perspective, where the perturbations of value estimate and gain in precision that occurs during deliberation endure and remain present for post-choice ratings. Finally, we calculated SoA, confidence, and RT in the same way as for the statistical artifact account, except that now confidence and RT were based on the perturbed true values.

To test whether our results might be explainable as mere statistical artifacts, we examined the different datasets (empirical, value refinement simulations, and statistical artifact simulations) and compared the results of several key analyses (Lee & Pezzulo, 2022). First, we measured the mean SoA across all trials. The empirical and value refinement datasets both showed a mean SoA of 2, whereas the statistical artifact dataset showed a mean SoA of 0. Next, we tested the predictive effect of post- versus pre-choice ratings. Specifically, we regressed choice consistency on both dV1 and dV2 in the same model. With the empirical data, dV2 had a much greater predictive effect than dV1 (0.80 versus 0.34). Moreover, when regressing consistency on dV1 and dV2 separately, the percentage of variance explained was greater when using dV2 (r -squared: 0.50 versus 0.45). This pattern held with the value refinement simulated data (beta weights for dV1 and dV2 in the same regression: 0.29 and 0.77; r -squared in separate regressions: 0.45 and 0.52). The pattern did not hold with the statistical artifact simulated data (beta weights for dV1 and dV2 in the same regression: 0.53 and 0.52; r -squared in separate regressions: 0.47 and 0.47). See Figure S1A-C.

Pre- and post-choice ratings are highly correlated, as they are intended to provide identical measures (apart from rating noise and potential choice-induced preference change). So, examining

regression weights when using both variables in the same model might not be advisable. We therefore conducted an additional analysis that allowed us to more properly compare the relative predictive power of pre- and post-choice ratings, as proposed by Lee and Pezzulo (Lee & Pezzulo, 2022). Specifically, we ran two similar regressions of consistency on dV1 and dV2: one in which dV1 is first regressed out of dV2 and one in which dV2 is first regressed out of dV1. This allowed us to test for how informative post-choice ratings were even after pre-choice ratings were considered and vice-versa. With the empirical data, in the different regressions, dV1 and dV2 (not the residuals) had approximately equal predictive effects (1.08 and 1.09, respectively). However, the incremental predictive effect of dV2 residuals was much greater than the incremental predictive effect of dV1 residuals (0.80 versus 0.34). This pattern held with the value refinement simulated data (dV1: 0.94; dV2: 1.02; dV1 residuals: 0.29; dV2 residuals: 0.77). The pattern did not hold with the statistical artifact simulated data (dV1: 0.97; dV2: 0.97; dV1 residuals: 0.53; dV2 residuals: 0.52). See Figure S1D-F.

If SoA captures meaningful preference change that occurs during choice deliberation and is instrumental to the developing choice, post-choice ratings should have a greater predictive effect than pre-choice ratings on choice confidence and RT. On the contrary, if SoA is purely a statistical artifact, the predictive effect of any rating session should be identical. To test this, we regressed RT and confidence, separately, on |dV1| and |dV2| simultaneously. With the empirical data, |dV2| had a greater predictive effect than |dV1| in both regressions (RT: beta for |dV1| = -0.10; beta for |dV2| = -0.24; confidence: beta for |dV1| = 0.14; beta for |dV2| = 0.32). This pattern held with the value refinement simulated data (RT: beta for |dV1| = -0.24; beta for |dV2| = -0.48; confidence: beta for |dV1| = 0.18; beta for |dV2| = 0.33). The pattern did not hold with the statistical artifact

898 simulated data (RT: beta for $|dV1|$ = -0.36; beta for $|dV2|$ = -0.37; confidence: beta for $|dV1|$ =
 899 0.25; beta for $|dV2|$ = 0.25). See Figure S1G-I.

901 *Figure S1.* Patterns of regression coefficients for the predictive effects of dV1 and dV2 on choice
 902 consistency, RT, and confidence are similar in the empirical data (left column) and the data
 903 generated by the value refinement account (center column), but fundamentally different in the data
 904 generated by the statistical artifact account (right column).
 905

906 In addition to the standard measure of SoA, we also report a measure of intra-choice SoA
 907 that we call SoA*. Because this measure has not been reported before, it was not examined in the
 908 study of Lee and Pezzulo (Lee & Pezzulo, 2022). Nonetheless, we followed the methods prescribed
 909 in that study as we compared the predictions of the value refinement and statistical artifact accounts
 910 against the empirical data. For this, we simulated stochastic choices rather than deterministic

choices, as this would maximize the possibility that the statistical artifact account might be able to explain the empirical data. Specifically, we simulated “choice ratings” under each of the accounts in exactly the same way as we did for the post-choice ratings. We then declared that the chosen option was the one with the higher choice rating and calculated SoA* accordingly. It turns out that both accounts generated a similar level of average SoA* (5.1 and 4.6). However, the statistical artifact account generated an average SoA of 0, whereas the value refinement account generated an average SoA of 1.3. We repeated this analysis using deterministic choices, and the statistical artifact account generated both SoA and SoA* of 0, whereas the value refinement account generated SoA of 1.6 and SoA* of 1.5.

Task Instructions

Exposure

“Before the evaluation part of the study begins, you will be presented with images of the different snacks, one at a time, in a random order. For each snack, you will briefly see a representative image in the center of the screen. You do not need to do anything during this stage, simply pay attention to which snacks are available in the set. The purpose of this stage is only to make you generally aware of the full set of snacks that you will be evaluating during the main part of the study.”

Rating

“You will now begin the value estimate stage of the study. Here, you will again be presented with images of snack foods (one at a time, in a random order). For each snack, you will see an image in the center of the screen. Below the image, you will see the question, ‘Do you like this snack?’ Below this question, you will see a slider scale where you can enter your rating

response. Click anywhere on the scale according to your personal preferences regarding eating this snack. The furthest left end of the scale indicates that you hate this snack. The furthest right end of the scale indicates that you love this snack. *Try to imagine committing to eating this snack on a frequent basis, not just one single time.* Once you have entered your response for a given snack, you will need to click the Enter button at the bottom of the screen to proceed.”

Standard choice

“You are about to begin the choice stage of the study. Here, you will be presented with pairs of images of snack foods (one pair at a time, in a random order). Below each pair of images, you will see the question, ‘Which snack do you prefer?’ Click on the image of the snack that you prefer. *Try to imagine committing to eating the snack that you choose on a frequent basis, not just one single time.*”

Implicit choice

“You are about to begin the choice stage of the study. Here, you will be presented with pairs of images of snack foods (one pair at a time, in a random order). Below each pair of images, you will see a slider scale along with the question, ‘Do you like this snack?’ Please enter your rating for how much you like each snack on the respective slider scale. The left-most end of the scale indicates that you hate the snack, the right-most end of the scale indicates that you love the snack; the center of the scale indicates that you are neutral about the snack. At the top of the screen, you will also see the question, ‘Which snack do you prefer?’ You will indicate your preferred snack via your ratings on the slider scales. For example, if you give a higher rating to the snack on the left than to the snack on the right, it indicates that you prefer the snack on the left. After you have entered both ratings on the slider scales, click the enter button to finalize your choice. *Try to*

956 *imagine committing to eating the snack that you choose on a frequent basis, not just one single*
957 *time.”*

958 *Implicit plus standard choice*

959 “You are about to begin the choice stage of the study. Here, you will be presented with
960 pairs of images of snack foods (one pair at a time, in a random order). Below each pair of images,
961 you will see a slider scale along with the question, ‘Do you like this snack?’ Please enter your
962 rating for how much you like each snack on the respective slider scale. The left-most end of the
963 scale indicates that you hate the snack, the right-most end of the scale indicates that you love the
964 snack; the center of the scale indicates that you are neutral about the snack. At the top of the screen,
965 you will also see the question, ‘Which snack do you prefer?’ Click on the image of the snack that
966 you prefer, *after* you enter both ratings on the slider scales. If you try clicking on one of the snacks
967 before you enter both ratings, an error message will pop up. *Try to imagine committing to eating*
968 *the snack that you choose on a frequent basis, not just one single time.”*