

A Self-directed Expanded Judgment Paradigm: Isolating the Pairwise Mechanism of the Attraction Effect

Tapas Rath (tapasr@iitk.ac.in)¹, Nisheeth Srivastava (nsrivast@iitk.ac.in)^{1,2} & Narayanan Srinivasan (nsrini@iitk.ac.in)¹

¹Department of Cognitive Science, Indian Institute of Technology Kanpur, Kanpur, India

²Department of Computer Science, Indian Institute of Technology Kanpur, Kanpur, India

Abstract

The attraction effect—where a decoy option increases preference for a dominating target—is a cornerstone of context-dependent choice, yet it is paradoxically fragile. Sequential accounts propose that context effects depend on which pairwise comparisons are emphasized during deliberation, but existing tests often confound comparison availability with memory/recency. We introduce a Self-directed Expanded Judgment paradigm in which participants repeatedly unblur and judge available option pairs, allowing information search to be observed and constrained. In Experiment 1, we validate the method by replicating a positive attraction effect. In Experiment 2, we causally manipulate comparison availability by disabling specific pairwise links while equalizing cumulative stimulus exposure at the decision stage (top-up control). Consistent with the preregistered hypothesis, the target’s relative share (RSTew) differed reliably between conditions, with the disabled Competitor-Decoy pair producing a positive attraction effect. In exploratory analyses, the disabled Target-Decoy pair produced negative attraction (repulsion). These findings provide causal evidence that the accessibility of specific pairwise comparisons can modulate—and under some conditions reverse—context effects.

Keywords: Context effects; Attraction effect; Pairwise comparison; Information search; Sequential decision making

Introduction

Human preferences are systematically shaped by context: the same option can become more or less attractive depending on the other options available. A canonical demonstration is the attraction effect, in which adding an option (the *decoy*) that is clearly worse than one of the focal alternatives (the *target*) increases the choice of that alternative. This phenomenon is theoretically important because it violates core rationality constraints such as *regularity* (Luce, 1977) and *independence from irrelevant alternatives*.

At the same time, the attraction effect is paradoxically fragile. Small changes to the task environment can attenuate the effect or even produce qualitatively different context effects. This fragility has motivated an influential mechanistic proposal: multi-alternative choices may not be evaluated holistically, but instead constructed through sequential pairwise comparisons that unfold over time (Noguchi & Stewart, 2014; Spektor et al., 2021; Trueblood et al., 2013). In this view, context effects depend on which comparisons are made salient or likely during deliberation, rather than being an inevitable consequence of the option set.

A key empirical challenge for testing pairwise-comparison accounts is separating limits on the opportunity to compare options from limits imposed by cognitive state. Many paradigms that manipulate the sequence of information confound these two factors: the same design choices that make certain comparisons more or less likely also change what must be retrieved from memory, how recently items were seen, or how evenly they were sampled. As a result, it is often unclear whether observed changes in context effects reflect structural constraints on which comparisons could be made, or downstream effects of memory decay, recency, or unequal exposure. This confound is particularly acute in sequential-presentation designs. For example, Evans et al. (2021) reversed the attraction effect by presenting options one at a time, requiring participants to compare currently visible items against memory representations of previously seen options. While this result highlights the role of sequential processing, it leaves open whether the reversal arose because certain pairwise comparisons were structurally unavailable, or because those comparisons were cognitively degraded by reliance on memory.

Furthermore, attempts to study context effects in other paradigms where decision-relevant information is revealed over time—including experience-based designs—have yielded inconsistent results: the attraction effect has been reported to attenuate, disappear, or even reverse when options must be learned from sampled outcomes (Ert & Lejarraaga, 2018; Hadar et al., 2018; Spektor et al., 2019), while others successfully replicated standard attraction effects even in experience-based settings (Balakrishnan et al., 2020). These mixed results suggest multiple confounds are at play. One key confound in experience-based designs is that participants must simultaneously pursue two distinct learning goals: learning the objective values of the options (through probability learning) and forming subjective preferences for different levels of those values (Liu & Trueblood, 2023). For example, when shopping for a new mattress, a person must simultaneously learn the objective properties—firmness levels—while discovering their own subjective preferences for different degrees of firmness. This dual learning problem represents another instance in which observed context effects conflate choice mechanisms with auxiliary cognitive demands.

We introduce a Self-directed Expanded Judgment paradigm designed to isolate pairwise mechanisms while minimizing these confounds. On each trial, participants are presented with a blurred three-option set (Target, Competitor, Decoy),

but can only unblur and inspect *two options at a time* via dedicated comparison controls. This interface externalizes information search: participants may re-view available pairs as often as they wish, allowing us to study how the *structure* of accessible comparisons shapes preference without forcing a fixed sequence. Critically, our paradigm treats the comparison structure itself as an experimental object: by selectively disabling specific pairwise comparisons, we can intervene on which relational evidence is even *possible* to obtain within a trial, rather than merely biasing attention toward particular options.

Our paradigm can be viewed as a self-directed instantiation of the expanded-judgment tradition, in which evidence is presented sequentially to make the evidence stream observable to both the decision-maker and the experimenter (Brown et al., 2009; Irwin et al., 1956). Unlike classic expanded-judgment tasks where the experimenter controls the sequence and timing of evidence samples—and unlike prior designs that presented options one at a time (Evans et al., 2021)—our interface presents options in explicit pairwise comparisons and allows the decision-maker to control both the order and duration of sampling. This converts the typically simultaneous option display into a decision-maker-controlled sequential process, while keeping critical comparisons perceptually available through explicit pairwise displays, thereby reducing the need to compare current inputs against memory traces. In this sense, our approach extends the expanded-judgment logic to multi-alternative choice by making the comparison process itself directly observable (Irwin et al., 1956; Malhotra et al., 2017; Tsetsos et al., 2012; Vickers et al., 1985).

We report two experiments using this paradigm. In Experiment 1, we validate the self-directed method by applying it to the preferential stimulus set of Noguchi and Stewart (2014) to confirm it recovers the basic attraction signature.

In Experiment 2, we test the causal role of pairwise availability. We manipulate the comparison structure within-subjects (disabling CD vs. TD links) while implementing a novel exposure equalization (‘top-up’) procedure to decouple comparison history from exposure duration. This allows us to isolate whether the attraction effect depends on the specific architecture of accessible comparisons.

Experiment 1

Introduction

The primary goal of Experiment 1 was to validate the Self-directed Expanded Judgment paradigm. Before using this novel interface to intervene on comparison availability (as planned for Experiment 2), it is necessary to establish that the paradigm itself does not disrupt the standard attraction effect. Previous attempts to study context effects in sequential or experience-based tasks have produced mixed results, often failing to replicate the effect or observing reversals due to probability learning confounds (Ert & Lejarraga, 2018).

Therefore, Experiment 1 asks a fundamental feasibility question: Can a self-directed, pairwise viewing interface elicit

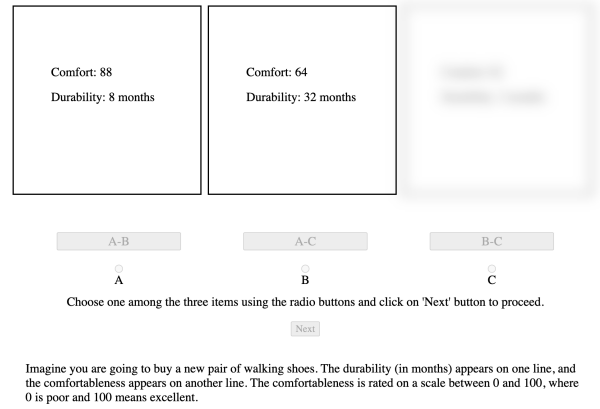


Figure 1: Experiment 1 Sample Trial




Figure 2: Experiment 2 Sample Trial

the standard attraction effect when using stimuli known to produce the effect in static description-based tasks? To test this, we utilized the preferential stimulus set from Noguchi and Stewart (2014), which has robustly demonstrated the attraction effect in traditional designs. Unlike Experiment 2, this validation experiment did not enforce statistical exposure equalization (the ‘top-up’ procedure) or restrict pairwise comparisons. Instead, participants were free to sample all pairwise relations naturally. We hypothesized that under these baseline conditions, the self-directed paradigm would successfully replicate the standard positive attraction effect ($RST > 0.5$).

Method

Participants Participants were recruited from the university, the ethics review committee of which approved the study. Participants received monetary compensation for their time. The final sample consisted of $N = 40$ participants. All participants had normal or corrected-to-normal vision and provided informed consent.

Stimuli We used the preferential stimulus set described by Noguchi and Stewart (2014). The stimuli were hypothetical consumer goods defined by two attributes (e.g., Quality and

Price). Refer to Figure 1 for a sample trial. Three options were generated for each trial: a Target (T), a Competitor (C), and a Decoy (D). The Target and Competitor formed the core choice set, trading off across the two attributes such that neither dominated the other. The Decoy was constructed to be asymmetrically dominated by the Target (i.e., worse than the Target on both attributes) but not by the Competitor. This structure is the standard definition of the attraction effect Huber et al. (1982). To prevent location effects, the spatial position of the options on the screen was randomized across trials.

Procedure The experiment was conducted in the lab using a custom web-based interface. On each trial, participants were presented with three blurred boxes representing the options. To sample information, participants clicked a designated comparison button to unblur a specific pair (e.g., Option A vs. Option B). Crucially, every sampling event required an immediate micro-judgment: participants were required to indicate which of the two visible items had the higher criterion value by clicking directly on that item. This action recorded the local judgment and immediately re-blurred the pair, returning the participant to the selection state to initiate the next comparison. This process was entirely self-directed: participants could view the pairs in any order and for any duration, and they could revisit pairs as many times as they wished. This feature served as an external memory store, reducing the cognitive load associated with maintaining attribute values in working memory. Once participants felt they had gathered sufficient information, they proceeded to the decision phase.

Prior to the main task, participants completed a 6-trial practice block to familiarize themselves with the interface. The practice trials mirrored the main experiment’s structure but used a simplified task (choosing the largest two-digit integer). Participants were required to answer at least 5 out of 6 trials correctly to proceed. The main experiment followed a within-subject design. Each participant completed 20 trials and took around 30 minutes on average.

Data Analysis We quantified context effects using the equal-weights version of the Relative Choice Share of the Target (RST_{EW}), a robust measure for triplet designs (Katsimpokis et al., 2022; Wedell, 1991). This metric averages the target’s relative advantage across the two counterbalanced decoy positions (favoring Option X vs. favoring Option Y), neutralizing any inherent preference for the core items. It is calculated as:

$$RST_{EW} = \frac{1}{2} \left(\frac{T_X}{T_X + C_X} + \frac{T_Y}{T_Y + C_Y} \right)$$

where T_k and C_k denote the counts of Target and Competitor choices when the decoy is positioned to favor option $k \in X, Y$. An RST_{EW} significantly greater than 0.5 indicates a decoy-induced bias towards the target.

Results

Experiment 1 produced a positive attraction effect. A one-sample two-tailed t -test was performed to compare RST val-

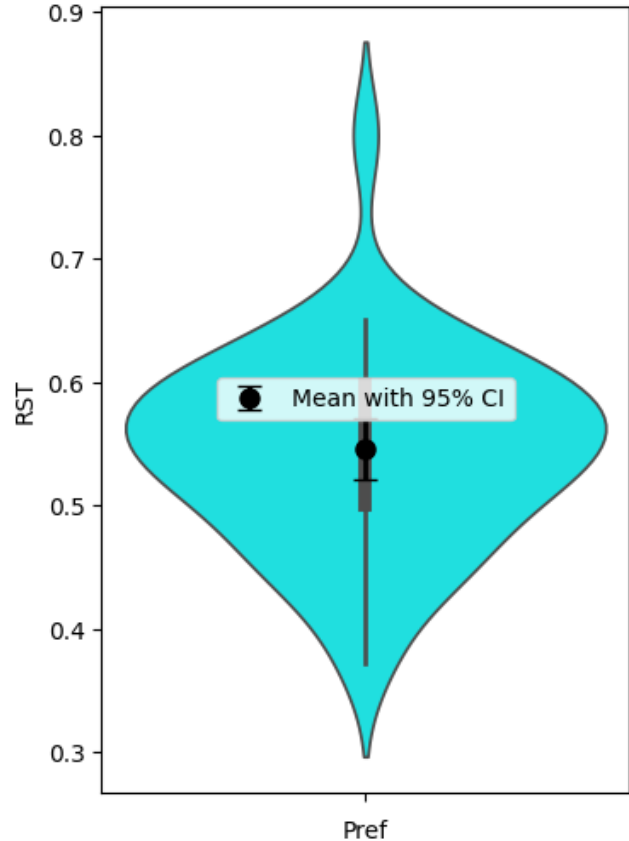


Figure 3: Exp1 RST distribution

ues against the null value of 0.5. The mean RST ($M = 0.546$, $SD = 0.079$) was significantly higher than the null value of 0.5, $t(39) = 3.648$, $p < 0.001$, Cohen’s $d = 0.577$. Figure 3 displays the violin plot of RST values along with 95% confidence intervals.

Experiment 2

Introduction

Experiment 1 established that the self-directed paradigm elicits the standard attraction effect. Experiment 2 now tests the causal mechanism: does the effect depend on the availability of specific comparisons?

To test this, we introduced two design changes. First, we used the commensurable “bar stimuli” from Spektor et al. (2022). As argued in recent work (Rath & Marupudi, 2025; Rath et al., 2025), these stimuli maximize the diagnostic value of the Competitor–Decoy (CD) relation, making its availability a critical theoretical lever. Second, we implemented an exposure top-up to ensure that any choice shifts are driven by relational information rather than unequal viewing time.

We preregistered a specific causal hypothesis on OSF (<https://osf.io/ve25x>): if attraction is driven by T-D dominance, then disabling the C-D comparison should preserve or enhance the effect (relative share of target $RST_{EW} > 0.5$), whereas

disabling the T-D comparison should diminish it. Whether the latter produces a full reversal (repulsion) was treated as exploratory.

Method

Participants Data collection followed a Bayesian sequential design (as preregistered). Sampling commenced with a minimum of 20 participants and proceeded in increments of 5 until a Bayes factor threshold was reached or a maximum of $N = 100$ was obtained. The final sample consisted of 20 participants recruited from the university, who received monetary compensation for participation. All participants had normal or corrected-to-normal vision.

Stimuli We used the commensurable “bar stimuli” set (Spektor et al., 2022). Each stimulus consisted of two orthogonally placed thin grey bars (one vertical, one horizontal) within a square boundary box. The criterion value of a stimulus was defined as the sum of the two bar lengths (max 200 units each; max total 400), and participants were instructed to choose the option with the highest total fill length (Refer to Figure 2 for a sample trial). Stimuli were generated algorithmically. Core pairs were sampled from uniform ranges ($h \in [160, 200]$, $v \in [70, 110]$).

Decoys were generated using a simulation-based procedure to ensure that the metric distance between Target and Decoy was strictly invariant to attribute orientation (i.e., matched between Wide-Low [WL] and Narrow-High [NH] targets). We simulated candidate decoy-generation policies—comparing fixed absolute versus relative attribute decrements—and evaluated the resulting distributions of target–decoy distance (Δ_{tot}) using Kolmogorov–Smirnov tests.

Based on these simulations, we selected a relative-decrement policy: Range and Frequency decoys were created by reducing the target’s manipulated attribute by a fixed percentage, while Range-Frequency decoys reduced both attributes by half that percentage (Huber et al., 1982). This policy minimized systematic deviations in decoy strength between WL and NH trials. The final stimulus set consisted of 24 core pairs (generated from a fixed seed) expanded into 72 main experimental trials, fully balanced for decoy type and target identity

Design The experiment employed a within-subject design with *Comparison Availability* as the manipulated factor. The trial types were CD-disabled (Target–Decoy enabled; Competitor–Decoy disabled), TD-disabled (Competitor–Decoy enabled; Target–Decoy disabled), and CT-disabled (Target–Competitor disabled, included for completeness). Each participant completed 84 trials: 72 experimental trials (24 per condition) and 12 catch trials where options had strictly ordered values (dominance > 20%) to assess attention.

Procedure The procedure followed the Self-directed Expanded Judgment paradigm. In the self-directed comparison phase, participants viewed three blurred options and sampled

information by clicking comparison buttons to unblur specific pairs. As in Experiment 1, every sampling event required a micro-judgment: to close a view and proceed, participants had to click on the larger of the two visible items. Crucially, on each trial, one comparison button was disabled (e.g., in a CD-disabled trial, the Competitor–Decoy button was inactive), restricting participants to sampling only the remaining two relational pairs. To correct for the exposure asymmetry induced by this restriction, the self-directed phase was followed by an exposure equalization (“top-up”) sequence. Participants clicked a “Show options (briefly)” button (refer to Figure 2), triggering a computer-controlled sequence where each individual option was unblurred alone in a random order. The duration of each presentation was calculated in real-time to perfectly equalize the cumulative viewing time of the Target, Competitor, and Decoy across the entire trial. Once equalization was complete, all options were re-blurred and the final selection radio buttons were enabled.

Let E_T , E_C , and E_D be the cumulative exposure times for each option during the self-directed phase, and let $E_{max} = \max(E_T, E_C, E_D)$. A brief sequential display presented each item i for a duration $R_i = R_{min} + (E_{max} - E_i)$. This ensured that at the moment of decision, the total exposure ($E_i + R_i$) was identical for the Target, Competitor, and Decoy ($E_{total} = E_{max} + R_{min}$). After the top-up sequence, when the radio buttons were enabled, participants made their final choice.

Participants were trained on this procedure via a 6-trial practice session. To ensure they understood the restricted comparison logic and the top-up sequence, the practice enforced the same constraints (disabled buttons, mandatory top-up) but used a simplified integer-comparison task. A passing score of 5/6 was required to enter the main experiment.

Data Analysis We used the same equal-weight Relative Share of the Target (RST_{EW}) measure defined in Experiment 1. Primary inference focused on the within-subject contrast between the *CD-disabled* and *TD-disabled* conditions.

Results

In enrollment order, the preregistered stopping threshold was crossed at the first scheduled interim look ($N = 20$), yielding a directional Bayes factor of $BF_{10^+} = 632.84$. All confirmatory analyses reported below therefore use this stopped sample of 20 participants.

Mean RST_{EW} was .572 in the CD-disabled condition ($SD = .123$), .420 in the TD-disabled condition ($SD = .083$), and .520 in the CT-disabled condition ($SD = .085$). The mean within-participant difference between CD-disabled and TD-disabled trials was .152.

The preregistered directional Bayesian paired-samples t -test provided strong evidence for the hypothesis that $RST_{EW,CD} > RST_{EW,TD}$, $BF_{10^+} = 632.84$ ($BF_{01^+} = 0.0016$). The corresponding preregistered one-tailed frequentist paired-samples t -test yielded converging evidence: the mean difference in RST_{EW} (CD – TD) was .152 ($SD = .155$, 95% CI [.080, .225]), $t(19) = 4.41$, $p < .001$, Cohen’s $d_z = 0.99$.

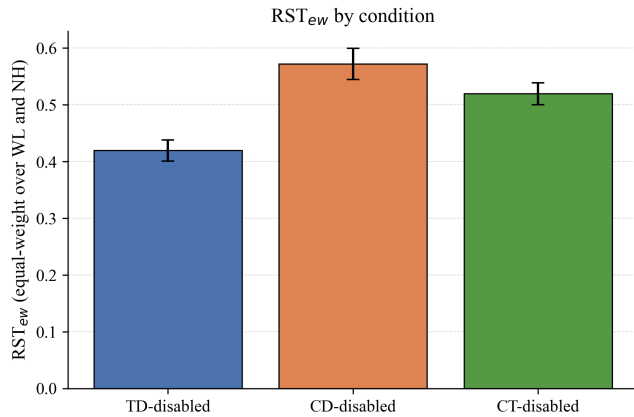


Figure 4: Exp2 RSTew by condition

As exploratory analyses, we examined whether RST_{EW} differed from chance (0.5) in each condition. Bayesian one-sample t -tests indicated strong evidence that RST_{EW} was below chance in the TD-disabled condition ($BF_{10} = 84.43$; $M = .420$, $SD = .083$, $N = 20$) and moderate evidence that it was above chance in the CD-disabled condition ($BF_{10} = 3.32$; $M = .572$, $SD = .123$, $N = 20$). For the CT-disabled condition, the evidence was anecdotal for no deviation from chance ($BF_{10} = 0.37$; $M = .520$, $SD = .085$, $N = 20$).

Frequentist one-sample t -tests showed the same general pattern. RST_{EW} was significantly below chance in the TD-disabled condition, $t(19) = -4.30$, $p < .001$, $d = -0.96$, and significantly above chance in the CD-disabled condition, $t(19) = 2.62$, $p = .017$, $d = 0.59$. No reliable deviation from chance was observed in the CT-disabled condition, $t(19) = 1.03$, $p = .318$, $d = 0.23$. Figure 4 displays RST_{EW} by condition with standard-error bars.

General Discussion

This research began with a paradox: the attraction effect is one of the most robust phenomena in decision science, yet it is notoriously fragile when task conditions change. We hypothesized that this fragility arises because multi-alternative choices are constructed through sequential pairwise comparisons, and that previous methods failed to isolate this mechanism from memory confounds. Our Self-directed Expanded Judgment paradigm was introduced to solve this, offering participants an “external memory store” and implementing a statistical “top-up” to equalize exposure.

The results provide strong support for the pairwise-comparison account of context effects. Experiment 1 demonstrated that a self-directed, pairwise viewing interface successfully elicits the standard attraction effect ($RST > 0.5$) using established preferential stimuli. This validates the paradigm as a legitimate tool for studying context effects, ruling out concerns that sequential sampling alone disrupts the effect (e.g., via probability learning confounds). Experiment 2 then showed that experimentally manipulating the availability of

pairwise comparisons while controlling for exposure yields a significant modulation of the effect. Consistent with our pre-registered hypothesis, the relative share of the target (RST_{EW}) was higher when the Target–Decoy comparison was available than when the Competitor–Decoy comparison was available. More critically, additional analyses indicated that this was not merely an attenuation: the CD-disabled condition produced a classic positive attraction effect, whereas the TD-disabled condition produced a negative attraction (repulsion) effect.

These findings offer a causal resolution to the debate over the fragility of context effects. By showing that we can turn the attraction effect “on” or reverse it into repulsion by toggling specific pairwise links, we demonstrate that the effect is not an intrinsic property of the option set. Instead, it is a conditional result of information search structure. This directly substantiates the hypothesis from the recent literature (Rath et al., 2025), which argued that the absence of attraction effects in perceptual tasks stems from the high diagnostic value of the Competitor–Decoy comparison. Here, we confirm that when this specific comparison is structurally disabled, the attraction effect re-emerges. Our results support sequential sampling models that rely on pairwise accumulation (Noguchi & Stewart, 2014; Rath et al., 2025; Trueblood et al., 2014). Under these accounts, preference is constructed by tallying binary wins. In a standard triplet, the Target accumulates ‘wins’ against the Decoy. However, if the Competitor also accumulates clear ‘wins’ against the Decoy (as in commensurable perceptual tasks), the Target’s advantage is neutralized. Our CD-disabled condition prevents the Competitor from accruing these wins, allowing the Target’s dominance to drive choice. Conversely, the emergence of repulsion when the Competitor–Decoy comparison is available suggests that when decision-makers focus on the relationship between these two options, the Competitor’s dominance over the Decoy becomes a salient cue that steers preference away from the Target.

Crucially, our design enables us to decouple the mechanism of attraction from common confounds. First, the self-directed pairwise interface externalizes the comparison process, ensuring that the effects are not driven by working memory failures or the difficulty of integrating sequential single-option samples (Evans et al., 2021). Second, because we equalized cumulative viewing time immediately prior to choice, the observed effects cannot be attributed to asymmetric exposure or simple attentional salience. The reversal from attraction to repulsion in the TD-disabled condition is therefore a structural consequence of the comparison graph itself.

One methodological limitation is the use of the top-up sequence. While necessary to control for exposure, it introduces a brief passive viewing phase that may dampen effect magnitudes compared to a purely active task. Future work should also examine the boundary conditions of these structural effects. Our use of commensurable bar stimuli in Experiment 2 was strategic—intended to maximize the salience of the CD relationship. We predict that the strong modulation observed here would diminish in domains with incommensurable at-

tributes (e.g., Quality vs. Price), where the Competitor–Decoy dominance relationship is inherently more ambiguous. Testing this prediction using a between-subjects manipulation of commensurability would provide a critical test of the pairwise-availability hypothesis.

Our paradigm opens a new avenue for interventionist decision science. Rather than inferring process from reaction times or eye-tracking, researchers can now actively design the comparison architecture of a choice environment. Future work can examine dynamic availability, such as comparisons that are available only for limited windows or come with a “cost,” to model interface constraints like scrolling and limited attention. Future work can also examine individual differences in search policy, testing whether some decision-makers preferentially seek CD comparisons while others preferentially seek TD comparisons, and how these policies relate to the sign and magnitude of context effects.

Conclusion

The attraction effect is neither a universal law nor a fragile artifact; it is a predictable output of the pairwise comparison process. By controlling how options are compared, we can control what is chosen. We provide causal evidence that isolating specific pairwise links—independent of memory or decision-stage salience—is sufficient to generate or reverse one of the most famous biases in behavioral science.

References

- Balakrishnan, H., Jagga, S., & Srivastava, N. (2020). Inducing preference reversals by manipulating revealed preferences. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 42.
- Brown, S., Steyvers, M., & Wagenmakers, E.-J. (2009). Observing evidence accumulation during multi-alternative decisions. *Journal of Mathematical Psychology*, 53(6), 453–462.
- Ert, E., & Lejarraga, T. (2018). The effect of experience on context-dependent decisions. *Journal of Behavioral Decision Making*, 31(4), 535–546.
- Evans, N. J., Holmes, W. R., Dasari, A., & Trueblood, J. S. (2021). The impact of presentation order on attraction and repulsion effects in decision-making. *Decision*, 8(1), 36.
- Hadar, L., Danziger, S., & Hertwig, R. (2018). The attraction effect in experience-based decisions. *Journal of Behavioral Decision Making*, 31(3), 461–468.
- Huber, J., Payne, J. W., & Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*.
- Irwin, F. W., Smith, W., & Mayfield, J. F. (1956). Tests of two theories of decision in an "expanded judgment" situation. *Journal of experimental psychology*, 51(4), 261.
- Katsimpokis, D., Fontanesi, L., & Rieskamp, J. (2022). A robust Bayesian test for identifying context effects in multi-attribute decision-making. *Psychonomic Bulletin & Review*. <https://doi.org/10.3758/s13423-022-02157-2>
- Liu, Y., & Trueblood, J. S. (2023). The effect of preference learning on context effects in multi-alternative, multi-attribute choice. *Cognition*, 233, 105365.
- Luce, R. D. (1977). The choice axiom after twenty years. *Journal of Mathematical Psychology*, 15(3), 215–233. [https://doi.org/10.1016/0022-2496\(77\)90032-3](https://doi.org/10.1016/0022-2496(77)90032-3)
- Malhotra, G., Leslie, D. S., Ludwig, C. J., & Bogacz, R. (2017). Overcoming indecision by changing the decision boundary. *Journal of Experimental Psychology: General*, 146(6), 776.
- Noguchi, T., & Stewart, N. (2014). In the attraction, compromise, and similarity effects, alternatives are repeatedly compared in pairs on single dimensions. *Cognition*, 132(1), 44–56.
- Rath, T., & Marupudi, V. (2025). Re-evaluating the numerical-perceptual distinction in the attraction effect. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 47.
- Rath, T., Srinivasan, N., & Srivastava, N. (2025). The attraction effect in perceptual decision-making: A case of dominance asymmetry. *Frontiers in Psychology*, 16, 1661748.
- Spektor, M. S., Bhatia, S., & Gluth, S. (2021). The elusiveness of context effects in decision making. *Trends in Cognitive Sciences*, 25(10), 843–854. <https://doi.org/10.1016/j.tics.2021.07.011>
- Spektor, M. S., Gluth, S., Fontanesi, L., & Rieskamp, J. (2019). How similarity between choice options affects decisions from experience: The accentuation-of-differences model. *Psychological review*, 126(1), 52.
- Spektor, M. S., Kellen, D., & Klauer, K. C. (2022). The repulsion effect in preferential choice and its relation to perceptual choice [Publisher: Elsevier]. *Cognition*, 225, 105164.
- Trueblood, J. S., Brown, S. D., & Heathcote, A. (2014). The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychological review*, 121(2), 179.
- Trueblood, J. S., Brown, S. D., Heathcote, A., & Busemeyer, J. R. (2013). Not just for consumers: Context effects are fundamental to decision making. *Psychological science*, 24(6), 901–908.
- Tsetsos, K., Chater, N., & Usher, M. (2012). Saliency driven value integration explains decision biases and preference reversal. *Proceedings of the national academy of sciences*, 109(24), 9659–9664.
- Vickers, D., Burt, J., Smith, P., & Brown, M. (1985). Experimental paradigms emphasising state or process limitations: I effects on speed-accuracy tradeoffs. *Acta Psychologica*, 59(2), 129–161.
- Wedell, D. H. (1991). Distinguishing among models of contextually induced preference reversals. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. <https://doi.org/10.1037/0278-7393.17.4.767>