*Judgment and Decision Making* (2026), **xx**:xx 1–12

CAMBRIDGE
UNIVERSITY PRESS

**Research Article**

# Pairwise Comparison Architecture in Context-Dependent Choice

Authors masked for review[1]

[1]Affiliations and contact information masked for review.

**Keywords:** attraction effect, context-dependent choice, pairwise comparison, expanded judgment, repulsion effect, process tracing

**Abstract**
The attraction effect is a central demonstration of context-dependent choice, yet it is fragile across presentation formats, stimulus domains, and process-tracing paradigms. We propose that this fragility partly reflects the architecture of comparison itself: context effects depend on which pairwise relations are available during deliberation, rather than arising as an inseparable property of the option set. To test this claim, we introduced a Self-Directed Expanded Judgment Paradigm in which participants revealed information by selecting pairwise comparison windows and making a brief micro-judgment before final choice. Experiment 1 validated the paradigm by showing that unrestricted self-directed pairwise sampling reproduced a positive attraction effect with established preferential stimuli. Experiment 2 then manipulated comparison availability directly. On each trial, one pairwise link was disabled while cumulative exposure to the target, competitor, and decoy was equalized before final choice. Disabling the Competitor–Decoy comparison produced attraction, disabling the Target–Decoy comparison produced repulsion, and disabling the Target–Competitor comparison produced no reliable deviation from chance. The preregistered contrast confirmed that target share was higher in the Competitor–Decoy-disabled condition than in the Target–Decoy-disabled condition. These findings provide causal evidence for a pairwise-comparison account in which attraction, repulsion, and null context effects are lawful outputs of the same comparison architecture under different structural constraints.

## 1. Introduction

The attraction effect, in which the choice share of a target option increases when an inferior decoy is added to a binary choice set, is one of the most theoretically consequential findings in judgment and decision making (Huber et al., 1982). Its significance lies in its violation of regularity and independence from irrelevant alternatives: adding an option should not increase the choice share of an existing alternative (Luce, 1977). Yet the effect is also notably fragile. Across perceptual, sequential, and experience-based paradigms, attraction can weaken, disappear, or reverse (Ert & Lejarraga, 2018; Evans et al., 2021; Hadar et al., 2018; Spektor et al., 2019, 2021, 2022). This fragility raises a mechanistic question: when a decoy changes choice, what process actually carries its influence?

One possibility is that the decoy alters the option set holistically, changing the valuation of the target and competitor as a whole. A different possibility is that preference is constructed from local relational verdicts, so that the decoy's influence depends on which specific pairwise comparisons are available during deliberation. On this pairwise-comparison view, the relevant evidence is not merely that a target, competitor, and decoy are present. What matters is whether the decision maker can compare Target with Decoy, Competitor with Decoy, and Target with Competitor, and how the outcomes of those comparisons accumulate before choice (Noguchi & Stewart, 2014; Trueblood et al., 2013, 2014).

Testing this claim requires more than observing which information participants sample. A key empirical challenge for pairwise-comparison accounts is separating limits on the opportunity to compare options from limits imposed by the cognitive state of the decision maker. 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 alter what must be retrieved from memory, how recently options were seen, and 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 consequences of memory decay, recency, or unequal exposure.

This problem is especially 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 options against memory representations of previously seen options. This result highlights the importance of sequential processing, but it leaves open whether the reversal arose because particular pairwise comparisons were structurally unavailable or because those comparisons were cognitively degraded by reliance on memory.

Related ambiguities arise in experience-based paradigms, where decision-relevant information is revealed through sampling. In such tasks, attraction effects have been found to attenuate, disappear, or reverse (Ert & Lejarraga, 2018; Hadar et al., 2018; Spektor et al., 2019). One reason these results are difficult to interpret is that participants must often solve two learning problems at once: learning the objective properties of the options and discovering their own subjective preferences over those properties (Liu & Trueblood, 2023). For example, when choosing a mattress, a person may need to learn both the available firmness levels and which degree of firmness they actually prefer. In such cases, observed context effects may reflect not only comparison architecture, but also auxiliary demands of memory, probability learning, and preference discovery.

To provide such a test, we developed a Self-Directed Expanded Judgment Paradigm (EJP). In this paradigm, participants initially see three blurred options. They acquire information by clicking buttons that reveal specific option pairs, and each reveal requires an immediate micro-judgment about which of the two visible options is superior. The procedure externalizes pairwise comparison while allowing participants to control the order, duration, and frequency of sampling. Critically, individual comparison buttons can be disabled, making it possible to manipulate which pairwise relations are structurally available rather than inferring comparison availability from search behavior after the fact.

The EJP paradigm extends the expanded-judgment tradition, in which the evidence stream is made observable to both the participant and the experimenter (Irwin et al., 1956). Unlike classic expanded-judgment tasks, where the experimenter controls the sequence of evidence samples, and unlike single-option sequential designs, the present interface presents information in explicit pairwise displays. This preserves the sequential character of deliberation while reducing the need to compare currently visible information against memory traces of previously seen options.

We report two experiments. Experiment 1 asked whether the EJP interface preserves the standard attraction effect under unrestricted comparison availability. This validation step was necessary because the interface is structurally novel and could, in principle, attenuate or distort context effects independently of any comparison-availability manipulation. Experiment 2 then delivered the critical causal test. Using commensurable bar stimuli, we selectively disabled one pairwise comparison link on each trial while equalizing cumulative exposure across options before final choice.

The pairwise-comparison account generated three qualitative predictions. When the Competitor–Decoy comparison is disabled, the Target's Target–Decoy advantage should remain unopposed, producing positive attraction. When the Target–Decoy comparison is disabled, the Competitor's advantage over the Decoy should dominate, producing repulsion. When the Target–Competitor comparison is disabled while both decoy-related comparisons remain available, context effects should remain near chance. This pattern would indicate that attraction and repulsion are not separate anomalies, but opposite outputs of the same comparison architecture under different structural constraints.

## 2. Method

### 2.1. General paradigm

Both experiments used the Self-Directed Expanded Judgment Paradigm to test whether the availability of specific pairwise comparisons modulates the attraction effect. On each trial, participants were presented with three blurred options and acquired information by clicking dedicated comparison buttons that revealed specific option pairs. Each reveal required an immediate micro-judgment indicating which of the two visible options had the higher criterion value. The pair was then re-blurred, and the display returned to the comparison-selection state. This procedure externalized the comparison process while reducing reliance on working memory, thereby making pairwise information search directly observable.

Across both experiments, context effects were quantified using the equal-weighted Relative Share of the Target ($\mathrm{RST}_{\mathrm{ew}}$), which averages the target's relative advantage across the two counterbalanced decoy positions and thereby adjusts for any baseline preference between the core options (Katsimpokis et al., 2022; Wedell, 1991). Values above .5 indicate a bias toward the target, values below .5 indicate a bias toward the competitor, and values at .5 indicate no bias.

### 2.2. Experiment 1

Experiment 1 served as a validation test of the EJP paradigm. Before using the paradigm to manipulate comparison availability, it was necessary to establish that the interface itself did not disrupt the standard attraction effect.

#### 2.2.1. Participants

Forty participants were recruited from the university student population, yielding a final analysis sample of $N = 40$. All participants had normal or corrected-to-normal vision, provided informed consent prior to participation, and received monetary compensation for their time. The experiment was approved by the university ethics review committee.

#### 2.2.2. Stimuli

Experiment 1 used the preferential stimulus set introduced by Noguchi and Stewart (2014). Stimuli consisted of hypothetical consumer goods defined by two attributes, such as quality and price. On each trial, three options were generated: a Target (T), a Competitor (C), and a Decoy (D). The Target and Competitor formed the core choice set and traded off across attributes such that neither dominated the other, whereas the Decoy was dominated by the Target. Spatial positions of the three options were randomized across trials.

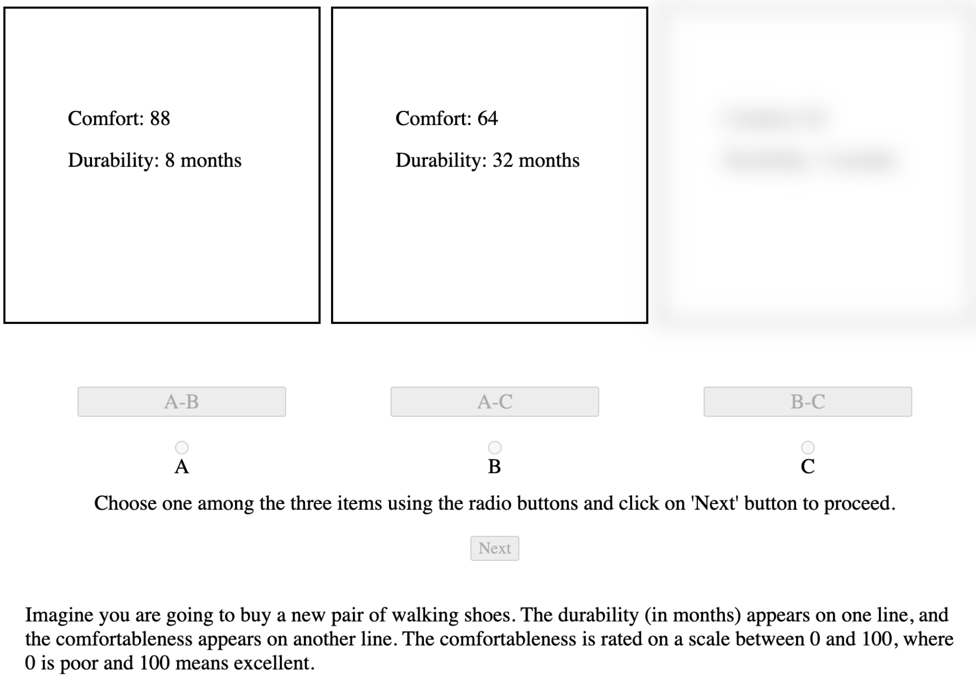


**Figure 1.** *Sample trial from the self-directed EJP validation task with all pairwise comparisons available..*

#### 2.2.3. Procedure

Participants could sample any available pairwise comparison in any order, for any duration, and as often as they wished. Each sampling event required a micro-judgment before the display re-blurred, after which participants could continue sampling or proceed to the final choice screen. Before the main task, participants completed a 6-trial practice block using a simplified integer-comparison task. A score of at least 5 out of 6 correct responses was required to proceed. The main experiment consisted of 20 trials and took approximately 30 minutes to complete.

### 2.3. Experiment 2

Experiment 2 used the validated EJP paradigm to test the causal role of pairwise comparison availability. Specifically, it examined whether disabling particular comparison links would systematically modulate the direction of the resulting context effect.

#### 2.3.1. Participants

Experiment 2 was preregistered prior to data collection on OSF (anonymized link for review). Data collection followed a preregistered Bayesian sequential sampling plan, beginning with a minimum of 20 participants and proceeding in increments of 5 until evidence for the directional alternative reached $\mathrm{BF}^{+}_{10} \geq 10$, evidence for the null reached $\mathrm{BF}^{+}_{01} \geq 10$, or the maximum sample size of $N_{\max} = 70$ was reached. Twenty participants were recruited from the university student population and received monetary compensation for participation. No participant was

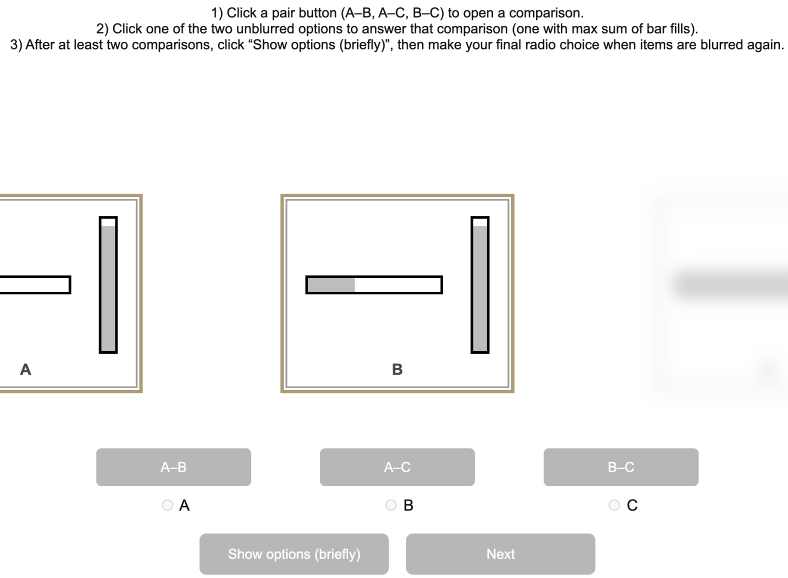


**Figure 2.** *Sample trial from the causal manipulation task illustrating commensurable bar stimuli, restricted comparison availability, and the exposure-equalization top-up stage..*

excluded from the stopped sample, yielding a final analysis sample of $N = 20$. All participants had normal or corrected-to-normal vision.

#### 2.3.2. Design

Experiment 2 employed a within-subject manipulation of comparison availability. On each trial, one comparison link was disabled, yielding three conditions: CD-disabled, TD-disabled, and CT-disabled. In the CD-disabled condition, the Competitor–Decoy comparison was unavailable; in the TD-disabled condition, the Target–Decoy comparison was unavailable; and in the CT-disabled condition, the Target–Competitor comparison was unavailable. Each participant completed 72 experimental trials, with 24 trials per condition, together with 12 catch trials in which the options were strictly ordered in value.

#### 2.3.3. Stimuli

We used commensurable bar stimuli adapted from Spektor et al. (2022). Each option consisted of two orthogonally oriented thin grey bars, one vertical and one horizontal, displayed within a square boundary. The criterion value of an option was defined as the sum of the two bar lengths, and participants were instructed to choose the option with the greatest total fill length (Figure 2). Core pairs were sampled from uniform ranges ($h \in [160, 200]$, $v \in [70, 110]$). Decoys were generated algorithmically so that target–decoy distance was matched across Wide–Low and Narrow–High target orientations. Full details of the simulation-based decoy-generation procedure and validation checks are reported in the Supplementary Materials.

#### 2.3.4. Procedure

The procedure followed the EJP paradigm established in Experiment 1, with two modifications. First, on each trial one comparison button was disabled, restricting participants to sampling

**Table 1.** *Experiment 2 condition-wise* $\mathrm{RST}_{\mathrm{ew}}$ *results..*

| Condition | $M$ | $SD$ | $BF_{10}$ | $t(19)$ | $p$ | $d$ |
|---|---|---|---|---|---|---|
| CD-disabled | .572 | .123 | 3.32 | 2.62 | .017 | .59 |
| TD-disabled | .420 | .083 | 84.43 | -4.30 | ¡.001 | -.96 |
| CT-disabled | .520 | .085 | .37 | 1.03 | .318 | .23 |

*Note.* Tests compare $\mathrm{RST}_{\mathrm{ew}}$ against .5. The primary preregistered analysis was the directional paired contrast between CD-disabled and TD-disabled conditions.

only the two remaining pairwise relations. Second, after the self-directed sampling phase, participants completed an exposure-equalization top-up sequence in which the three options were briefly presented individually so that cumulative exposure to the Target, Competitor, and Decoy was matched before final choice. Once equalization was complete, the final choice controls were enabled. Full implementation details of the equalization algorithm are reported in the Supplementary Materials.

### 2.4. Analysis plan

For Experiment 1, the primary question was whether $\mathrm{RST}_{\mathrm{ew}}$ exceeded .5, indicating that the EJP paradigm successfully reproduced the standard attraction effect. For Experiment 2, the preregistered primary analysis focused on the within-subject contrast between the CD-disabled and TD-disabled conditions. The central prediction was that $\mathrm{RST}_{\mathrm{ew}}$ would be higher in the CD-disabled condition than in the TD-disabled condition. The CT-disabled condition was included as an additional diagnostic control, and condition-wise tests against .5 were treated as secondary analyses. Bayes-factor $t$-tests used a Cauchy prior on standardized effect size with $r = .707$. All confirmatory tests treated participant as the unit of analysis; stimulus items were fixed by design and counterbalanced across target orientation and decoy type.

## 3. Results

### 3.1. Experiment 1

Experiment 1 produced the predicted positive attraction effect. A one-sample two-tailed $t$-test comparing $\mathrm{RST}_{\mathrm{ew}}$ against the null value of .5 showed that mean $\mathrm{RST}_{\mathrm{ew}}$ was significantly above chance ($M = .546$, $SD = .079$, $t(39) = 3.65$, $p < .001$, $d = .58$). This validates the EJP paradigm as a method that can recover the standard attraction effect under free-sampling conditions. Figure 3 shows the distribution of participant-level $\mathrm{RST}_{\mathrm{ew}}$ values with 95% confidence intervals.

### 3.2. Experiment 2

The preregistered stopping threshold was crossed at the first scheduled interim look ($N = 20$). All confirmatory analyses reported below use this stopped sample.

Descriptively, mean $\mathrm{RST}_{\mathrm{ew}}$ 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 the CD-disabled and TD-disabled conditions was .152. Table 1 summarizes the condition-wise analyses.

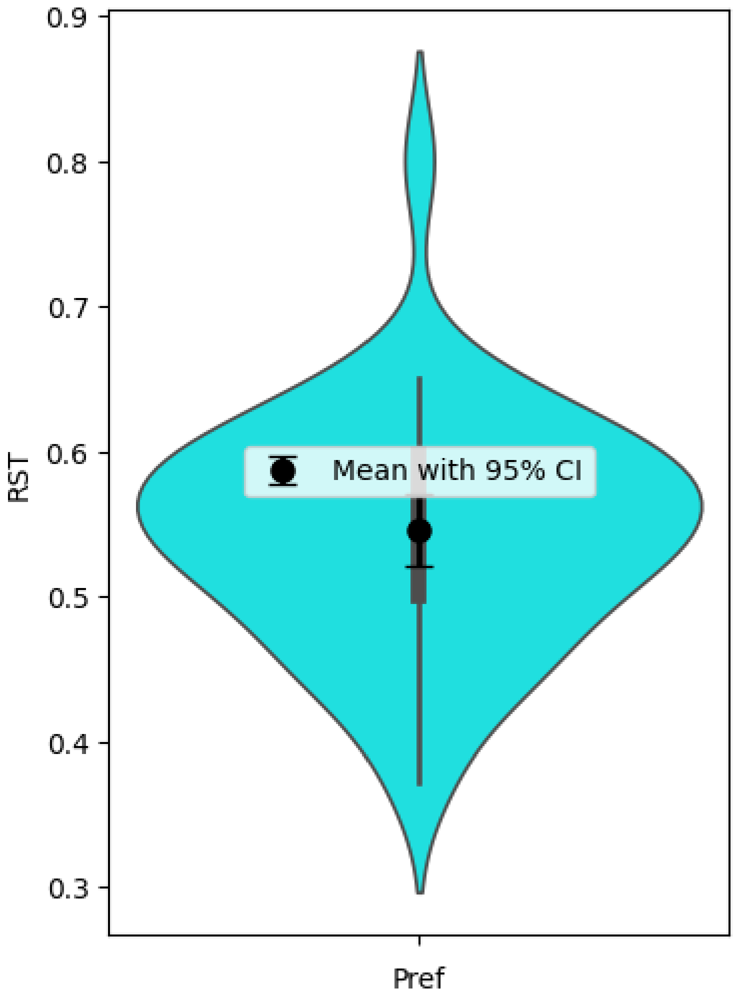


**Figure 3.** *Distribution of participant-level* $\mathrm{RST}_{\mathrm{ew}}$ *values in Experiment 1. Error bars represent 95% confidence intervals..*

The preregistered directional Bayesian paired-samples $t$-test provided strong evidence for the hypothesis that $\mathrm{RST}_{\mathrm{ew,CD}} > \mathrm{RST}_{\mathrm{ew,TD}}$ ($\mathrm{BF}_{10}^{+} = 632.84$, $\mathrm{BF}_{01}^{+} = .0016$). A corresponding one-tailed paired-samples $t$-test yielded converging evidence: the mean difference in $\mathrm{RST}_{\mathrm{ew}}$ (CD minus TD) was .152 ($SD = .155$, 95% CI [.080, .225], $t(19) = 4.41$, $p < .001$, $d_z = .99$).

As exploratory analyses, we tested whether $\mathrm{RST}_{\mathrm{ew}}$ differed from chance within each condition. Bayesian one-sample $t$-tests indicated strong evidence that $\mathrm{RST}_{\mathrm{ew}}$ was below chance in the TD-disabled condition ($BF_{10} = 84.43$), moderate evidence that it was above chance in the CD-disabled condition ($BF_{10} = 3.32$), and anecdotal evidence for no deviation from chance in the CT-disabled condition ($BF_{10} = .37$). Frequentist one-sample $t$-tests showed the same overall pattern: $\mathrm{RST}_{\mathrm{ew}}$ was significantly below chance in the TD-disabled condition ($t(19) = -4.30$, $p < .001$, $d = -.96$), significantly above chance in the CD-disabled condition ($t(19) = 2.62$, $p = .017$, $d = .59$), and did not differ reliably from chance in the CT-disabled condition ($t(19) = 1.03$, $p = .318$, $d = .23$). Figure 4 shows $\mathrm{RST}_{\mathrm{ew}}$ by condition with standard-error bars.

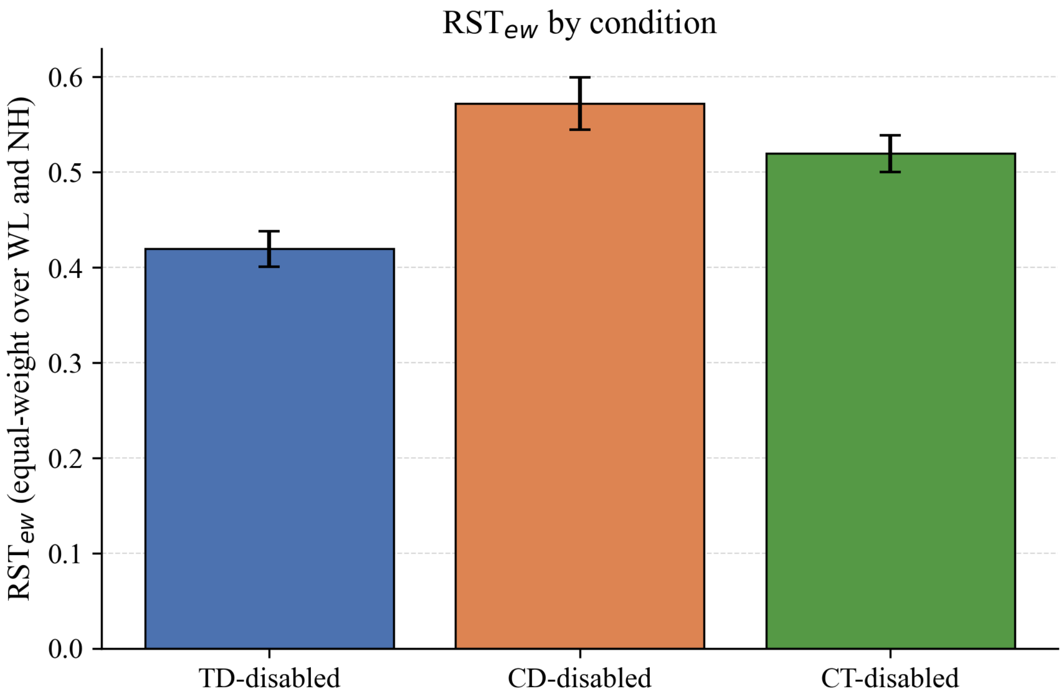


**Figure 4.** RST$_{\mathrm{ew}}$ *by comparison-availability condition in Experiment 2. Error bars represent standard errors..*

## 4. Discussion

Across two experiments, the Self-Directed Expanded Judgment Paradigm provided converging evidence that context effects depend on the structure of accessible pairwise comparisons rather than arising inevitably from the option set. Experiment 1 established the necessary methodological precondition by showing that the paradigm reproduces a standard positive attraction effect when all pairwise relations remain available and participants are free to sample them in a self-directed manner. This validation is important because prior sequential and experience-based paradigms have often yielded attenuated, absent, or reversed context effects, making it difficult to determine whether such outcomes reflect genuine changes in comparison structure or are instead driven by memory, recency, or learning confounds (Ert & Lejarraga, 2018; Evans et al., 2021; Spektor et al., 2022). By recovering the classic attraction pattern with established preferential stimuli (Noguchi & Stewart, 2014), the present findings show that the EJP interface itself does not abolish the effect.

Experiment 2 then used this validated paradigm to test the causal role of comparison availability directly. The preregistered contrast showed that target share was higher when the Competitor–Decoy comparison was disabled than when the Target–Decoy comparison was disabled, confirming the central prediction derived from the pairwise-comparison account. More importantly, the broader pattern across conditions was theoretically diagnostic: disabling the Competitor–Decoy link produced a positive attraction effect, disabling the Target–Decoy link produced repulsion, and disabling the Target–Competitor link left performance near chance. This pattern is difficult to reconcile with views that treat context effects as inseparable properties of the option set as a whole. Instead, it is more consistent with sequential models in which preference is constructed through the accumulation of outcomes from specific pairwise contests (Noguchi & Stewart, 2014; Trueblood et al., 2013, 2014).

On this interpretation, the decoy does not globally distort preference by its mere presence. Rather, its influence depends on which relational comparisons remain available during deliberation. When the Competitor–Decoy comparison is blocked, the competitor can no longer accrue the clear decoy-related wins that would otherwise offset the target's asymmetric advantage, allowing the target's dominance over the decoy to drive a positive attraction effect. When the Target–Decoy comparison is blocked, that privileged source of evidence is removed while the competitor's dominance over the decoy remains available, shifting preference away from the target and producing repulsion. The near-chance outcome in the Target–Competitor-disabled condition is especially informative: it is not merely a null result, but a structural constraint on the mechanism, indicating that decoy leverage is carried primarily through the Target–Decoy and Competitor–Decoy relations rather than through the direct Target–Competitor comparison.

These findings also speak to recent arguments about why attraction effects are often weak or absent in commensurable perceptual tasks. In such tasks, the Competitor–Decoy relation can become especially transparent, allowing the competitor to accumulate decoy-related wins that neutralize the target's advantage (Rath et al., 2025; Spektor et al., 2022). The present results provide a causal complement to that claim. Using commensurable bar stimuli adapted from Spektor et al. (2022), we show that when this specific comparison is structurally removed, a positive attraction effect re-emerges, whereas preserving it while removing the Target–Decoy link can reverse the effect altogether. Attraction and repulsion therefore appear not as unrelated anomalies, but as opposite outputs of the same underlying comparison architecture.

A central strength of the present design is that it isolates comparison availability from several confounds that have complicated earlier work. The EJP interface externalizes the comparison process by allowing participants to revisit available pairs rather than requiring them to compare currently visible options against degraded memory traces, extending the expanded-judgment logic to multi-alternative choice (Brown et al., 2009; Irwin et al., 1956). In Experiment 2, the exposure-equalization top-up further reduced the possibility that observed effects were driven simply by unequal cumulative viewing time or passive attentional salience. As a result, the shift from attraction to repulsion is more plausibly interpreted as a structural consequence of the comparison graph itself than as a by-product of exposure imbalance or recency.

Several limitations should be acknowledged. First, Experiment 2 used a relatively small sample, although the preregistered Bayesian stopping threshold was crossed at the first scheduled interim look and the primary contrast produced strong evidence. Replication with a larger and more diverse sample would strengthen confidence, especially for the condition-wise repulsion result in the TD-disabled condition. Second, the top-up procedure, while necessary to control exposure, introduced a brief passive viewing stage that may have dampened effect magnitudes relative to a purely self-directed task. Third, Experiment 2 intentionally used highly commensurable bar stimuli because these make the decoy-related comparisons maximally diagnostic (Rath et al., 2025; Spektor et al., 2022). Whether equally strong manipulations would be observed in domains with more ambiguous or incommensurable attributes remains an open question.

Future work should examine how comparison-availability effects vary across levels of commensurability, and should test more dynamic interventions in which specific comparisons become costly, transiently available, or sequence-dependent (Spektor et al., 2021). A second priority is to embed the present proposal in an explicit formal model that combines comparison selection with pairwise evidence accumulation. Such a model would provide a stronger bridge between the present causal manipulation and broader accumulator-style theories of multi-alternative choice (Brown et al., 2009; Noguchi & Stewart, 2014; Trueblood et al., 2014).

In conclusion, the present paper argues that the fragility of the attraction effect is not well understood if one asks only whether a decoy is present. A more informative question is which pairwise comparisons remain available while the decision maker constructs a preference. Across two experiments, the results support the view that attraction, repulsion, and null context effects are lawful outcomes of a structured comparison process.

**Acknowledgments.** None.

**Funding Statement.** This research received no specific grant from any funding agency, commercial, or not-for-profit sectors.

**Competing Interests.** The authors declare no competing interests.

**Data Availability Statement.** Data, stimulus materials, preregistration documents, analysis code, figure-generation scripts, and a codebook are available to reviewers through the anonymized OSF repository: https://osf.io/62wfr/overview?view_only=91f7f5edf2ce4684b5ee1b470ba30912. A public OSF link will be provided on acceptance.

**Ethical Standards.** The research was approved by the university ethics review committee, and all participants provided informed consent before participation.

**Author Contributions.** Author contribution details are masked for review.

## Supplementary Materials

### Supplementary Method S1: Simulation-Based Decoy Generation for Experiment 2

Experiment 2 used a commensurable bar-stimulus set derived from Spektor et al. (2022), in which each option consisted of two orthogonally oriented grey bars and the criterion value was defined as the sum of the horizontal and vertical bar lengths. Core option pairs were generated such that the two focal options traded off across dimensions while remaining comparable in total criterion value.

Core stimuli were constructed as Wide–Low (WL) options by sampling horizontal and vertical bar lengths independently from uniform ranges (horizontal: 160–200 units; vertical: 70–110 units). Corresponding Narrow–High (NH) options were generated by reversing the attribute values of each WL option. For each core pair, both WL-target and NH-target trials were simulated to ensure symmetry across target orientations.

A key design requirement was that decoy strength should not vary systematically as a function of target orientation. To achieve this, decoy generation rules were selected using a simulation-based policy comparison procedure. Candidate policies differed in how the decoy was constructed from the target, including fixed absolute decrements, proportional decrements, and total-criterion-based adjustments.

For each policy, decoy strength was quantified using relative distance measures. Specifically, we computed (i) the relative total distance between target and decoy, $(T_{\mathrm{tot}} - D_{\mathrm{tot}})/T_{\mathrm{tot}}$, and (ii) relative attribute-wise distances, $(T_h - D_h)/T_h$ and $(T_v - D_v)/T_v$. These measures capture the proportional deviation of the decoy from the target in overall and attribute-specific terms.

The distributions of these relative distance measures were compared between WL-target and NH-target trials using two-sample Kolmogorov–Smirnov tests. Policies were evaluated based on the KS statistic $D$, with lower values indicating greater invariance of decoy strength across target orientations.

Simulation results revealed that fixed absolute decrement policies produced substantial asymmetries across orientations, particularly for attribute-wise distances (KS $D \approx .60$), reflecting the fact that identical absolute changes correspond to different proportional changes across attribute scales. In contrast, proportional decrement policies yielded substantially smaller KS statistics (e.g., $D \approx .25$ for attribute-wise distances), indicating improved alignment of decoy strength across WL and NH orientations. The total-criterion-based rule showed intermediate performance.

Based on these results, the proportional decrement policy was selected for stimulus construction. Under this policy, Range and Frequency decoys were generated by reducing the relevant attributes of the target by a fixed percentage, while Range–Frequency decoys were generated by proportionally reducing both attributes, consistent with the logic of asymmetric-dominance manipulations (Huber et al., 1982).

The final stimulus set comprised 24 core pairs generated using a fixed random seed and expanded into 72 experimental trials, fully balanced across decoy type and target orientation. This simulation-based construction

**Table 2.** *Kolmogorov–Smirnov tests comparing distributions of relative distance measures between Wide–Low and Narrow–High target orientations across decoy-generation policies..*

| Rule | Total | Horizontal | Vertical |
|---|---|---|---|
| Absolute | .00 (1.000) | .60 (.001) | .60 (.001) |
| Proportional | .10 (1.000) | .25 (.571) | .25 (.571) |
| Total-based | .05 (1.000) | .30 (.336) | .25 (.571) |

*Note.* Values represent KS statistics with corresponding $p$-values in parentheses. Lower KS values indicate greater similarity between Wide–Low and Narrow–High distributions.

ensured that any observed differences in Experiment 2 could be attributed to the manipulation of comparison availability rather than unintended variation in decoy strength across stimulus orientations. Full KS statistics for all candidate policies and metrics are reported in Table 2.

## Supplementary Method S2: Exposure-Equalization Top-Up Procedure

Because Experiment 2 disabled one pairwise comparison link on each trial, the self-directed sampling phase would otherwise have produced a systematic exposure imbalance across the three options. Specifically, the option that appeared in both available comparisons would receive more cumulative viewing time than the two options connected by the disabled link. To prevent this asymmetry from confounding the interpretation of comparison-availability effects, the self-directed phase was followed by a computer-controlled exposure-equalization top-up procedure.

Let $E_T$, $E_C$, and $E_D$ denote the cumulative exposure times accrued by the Target, Competitor, and Decoy, respectively, during the self-directed sampling phase. Define the maximum exposure across the three options as

$$E^* = \max(E_T, E_C, E_D).$$

After participants completed self-directed sampling, they clicked a *Show options briefly* button that initiated the top-up sequence. During this sequence, each option $i \in \{T, C, D\}$ was presented individually, in randomized order, for an additional duration

$$R_i = R + (E^* - E_i),$$

where $R$ denotes a brief minimum display duration applied to all options. After the top-up, total cumulative exposure was therefore identical for all three options:

$$E_i + R_i = E^* + R \qquad \text{for all } i \in \{T, C, D\}.$$

This procedure was introduced to isolate the structural effect of comparison availability from any downstream influence of unequal cumulative viewing time at the moment of final choice. In other words, the top-up ensured that the disabled-link manipulation changed which relational evidence could be acquired, but did not leave one option more visually familiar merely because it had appeared in more self-directed sampling events. The final choice radio buttons were enabled only after the equalization sequence had been completed, and all options were re-blurred before choice. Participants were trained on this procedure in practice trials that enforced the same restricted-comparison and mandatory top-up structure as the main task.

Conceptually, the top-up was designed as a conservative control. It removes a simple exposure-based explanation of condition differences, allowing any observed shift in target share to be interpreted more specifically in terms of comparison structure. At the same time, because the procedure introduces a brief passive viewing phase after active sampling, it may reduce the magnitude of the comparison-availability effect relative to a purely self-directed paradigm. For that reason, the top-up should be understood as a control procedure aimed at strengthening causal interpretation rather than maximizing effect size.

## References


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. https://doi.org/10.1002/bdm.2064

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. A. S., & Mayfield, J. F. (1956). Tests of two theories of decision in an "expanded judgment" situation. *Journal of Experimental Psychology*, *51*(4), 261–268. https://doi.org/10.1037/h0041911

Katsimpokis, D., Fontanesi, L., & Rieskamp, J. (2022). A robust Bayesian test for identifying context effects in multiattribute 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. https://doi.org/10.1016/j.cognition.2022.105365

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

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., 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.

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