Judgment and Decision Making (2026), **XX**:XX 1–13

CAMBRIDGE
UNIVERSITY PRESS

Empirical Article

When Difficult Trade-Offs Make Decoys Work: Subjective Equivalence Geometry in the Attraction Effect

Author names masked for review[1]

[1]Department, Institution masked for review, City, Country.

Keywords: attraction effect, decoy effect, context effects, multiattribute choice, trade-off difficulty, subjective equivalence, decision making

Abstract

The attraction effect, in which an inferior decoy increases choice of the option that dominates it, is central to context-dependent choice but fragile across stimulus constructions. One explanation is that difficult inter-attribute trade-offs make decision makers rely on local dominance relations. We argue that this explanation is underspecified unless the structural requirements for attraction are separated from their psychological realization: absence of pre-decoy bias must correspond to subjective ambiguity between the core options, and asymmetric decoy dominance must correspond to an asymmetry in ease or clarity in decoy-related comparisons. Experiment 1 established a robust attraction effect in a mixed perceptual–numerical task in which participants chose the disc with the lowest price per unit area. Experiment 2 used a preregistered adjustment task to measure subjective equivalence geometry in circle-price and bar stimuli. Circle-price stimuli produced broader equivalence dispersion than bars and, critically, placed both the target offset and target-referenced decoy location inside participants' equivalence bands more often than bars. These findings show that trade-off difficulty can bundle the conditions that make attraction possible, explaining why positive-versus-null stimulus contrasts are informative but not by themselves causally diagnostic.

1. Introduction

Adding an inferior option can make a neighboring option more attractive. In the classic attraction effect, a decoy that is asymmetrically dominated by one option increases choice of the option that dominates it (Huber et al., 1982; R. D. Luce, 1977). The effect has been central to theories of context-dependent choice because it violates regularity and independence from irrelevant alternatives. It is also empirically fragile: modest changes in stimulus construction, display format, or task demands can attenuate or remove the effect (Rath et al., 2025; Spektor et al., 2018, 2021, 2022).

A common explanation appeals to trade-off difficulty. If people are unsure whether an advantage on one attribute compensates for a disadvantage on another, a decoy can provide a local comparative reason for favoring one option (Hedgcock & Rao, 2009; Huber et al., 2014; M. F. Luce et al., 2001; Simonson & Tversky, 1992; Tversky & Simonson, 1993). This explanation is related to attribute commensurability. Trade-off difficulty concerns uncertainty about whether gains on one dimension compensate for losses on another. Commensurability concerns whether different attributes can be translated into a common evaluative currency. When such a currency is available, decision makers can often integrate attributes directly; when no obvious exchange rate is available, the trade-off is harder to resolve (Banerjee et al., 2024; Chang, 2013; Hayes et al., 2025; Walasek & Brown, 2024). Even formally commensurable attributes may remain psychologically difficult when the relevant values are noisy, nonlinear, or hard to compute quickly. Trade-off difficulty is therefore best treated as a graded property of the implemented task, not as the mere presence or absence of a formal common unit.

This graded view matters because trade-off difficulty does not by itself name the proximal conditions needed for attraction. One general requirement is weak pre-decoy bias between the two core options. If the decision maker already clearly prefers one core option, a decoy has little opportunity to redirect choice; attraction is most interpretable when the Target–Competitor relation is close or uncertain (Huber et al., 2014; Pettibone, 2012). Recent analyses of triplet–triplet attraction-effect measures make the same point methodologically: apparent decoy effects are difficult to interpret when the two core options are not well balanced before the decoy is introduced (Rath et al., 2026). Weak pre-decoy bias is therefore not only a stimulus-design convenience; it is an inferential requirement for attributing a choice shift to the decoy. Second, work on perceptual attraction effects shows that nominal attribute-level dominance is not sufficient for attraction (Rath et al., 2025). A decoy can be inferior to the target in attribute space but fail to create a realized item-level advantage if participants integrate the full stimuli into a common criterion. Pairwise Target–Decoy and Competitor–Decoy choices are therefore useful diagnostics of whether the decoy-side asymmetry is reflected in choice.

The present paper asks that upstream question. It shifts the analysis from behavioral diagnosis to the subjective stimulus relations and psychophysical comparison geometry that produce those diagnostic patterns. The structural requirement of weak core-option bias corresponds to a psychological state in which the target and competitor are treated as effectively equivalent. We refer to this as core-option ambiguity, but the term is intentionally broad: ambiguity may reflect genuine indifference, equal liking, equal dislike, closeness in subjective value, uncertainty about the exchange rate between attributes, or difficulty comparing incommensurate attributes. Similarly, the structural requirement of item-wise asymmetric dominance corresponds to a psychological asymmetry in pairwise comparison: the Target–Decoy comparison must be easier, clearer, or more decisive than the Competitor–Decoy comparison.

The central claim is that inter-attribute trade-off difficulty can produce both psychological conditions at once. When attributes are hard to integrate, subjective equivalence regions may broaden. That broadening can make the target and competitor ambiguous, satisfying A, while also making the competitor difficult to distinguish from the decoy. If the target remains locally and clearly superior to the decoy, this same geometry can satisfy S. Thus, trade-off difficulty can jointly toggle both inputs of the attraction effect. Experiment 1 establishes attraction in an incommensurate mixed perceptual–numerical task. Experiment 2 measures the underlying equivalence geometry and shows how the two inputs can become bundled.

This creates an identification problem. A contrast between difficult and easy stimulus families may compare a bundled $A = 1, S = 1$ case against a bundled $A = 0, S = 0$ case, where A denotes core-option ambiguity and S denotes decoy-related comparative asymmetry. Such diagonal evidence can show that some stimulus families produce attraction and others do not, but it cannot reveal which ingredient is operative. Table 1 summarizes this logic.

We report two experiments. Experiment 1 asks whether a mixed perceptual–numerical stimulus family can produce attraction. Participants chose among circular discs that varied in area and price, selecting the option with the lowest price per unit area. Experiment 2 then asks why such a result is not mechanistically diagnostic. Using a preregistered adjustment paradigm, we measured the subjective equivalence geometry of the circle-price family and of a near-commensurate bar family. The central prediction was that circle-price stimuli would produce broader equivalence regions and would more often include both the target offset and target-referenced decoy location inside those regions.

Table 1. *Design Space of Core-Option Ambiguity (A) and Decoy-Related Comparative Asymmetry (S).*

	$S = 0$	$S = 1$
$A = 0$	Null or weak context dependence	Theoretically degenerate
	Easy or commensurate stimuli; bar-like cases	Strong core preference leaves little room for decoy influence
$A = 1$	Critical off-diagonal case	Attraction predicted
	Requires a separate causal dissociation	Circle-price stimuli in the present studies

Note. The present paper tests whether common stimulus contrasts can occupy the diagonal cells by jointly moving A and S. It does not claim to test the off-diagonal case directly.

2. Experiment 1: Mixed Perceptual–Numerical Choice

Experiment 1 tested whether attraction emerges in a task that requires integration of a perceptual magnitude with a symbolic numerical magnitude. Each option combined circular disc area with price, and participants selected the disc with the lowest price per unit area. If inter-attribute trade-off difficulty supports attraction, this stimulus family should produce a reliable shift toward the target when a decoy is present.

2.1. Method

2.1.1. Participants

Thirty-six participants ($M_{\text{age}} = 21.11$ years, $SD = 3.42$; 27 male, 9 female) with normal or corrected-to-normal vision took part after providing informed consent. Six participants were excluded for failing catch trials (accuracy $< .80$), leaving a final sample of $N = 30$.

2.1.2. Stimuli and Procedure

Each trial displayed three black-filled circular discs on a white background, arranged in a triangular configuration with vertical positions jittered across trials. Each disc represented a hypothetical alloy, with its price displayed at the center of the disc. The decision-relevant attributes were perceptual disc area and numerical price (Figure 1).

Stimuli were generated following procedures adapted from Trueblood et al. (2013). One set of discs was sampled from a bivariate normal distribution with mean area 28,000 pixels and mean price 28 units. A second set was matched in price per unit area but had twice the area of the first set. These two sets served as the target and competitor, counterbalanced across trials. Decoys were positioned closer to one of the two core options in attribute space, creating two contextual conditions. Range, frequency, and range-frequency decoys were included (Huber et al., 1982).

Participants completed 36 trials in each contextual condition, for 72 experimental trials, plus 12 catch trials in which one option clearly had the lowest price per unit area. The 84 trials were randomized. Participants chose using the left, up, or right arrow keys and were instructed to respond quickly and accurately.

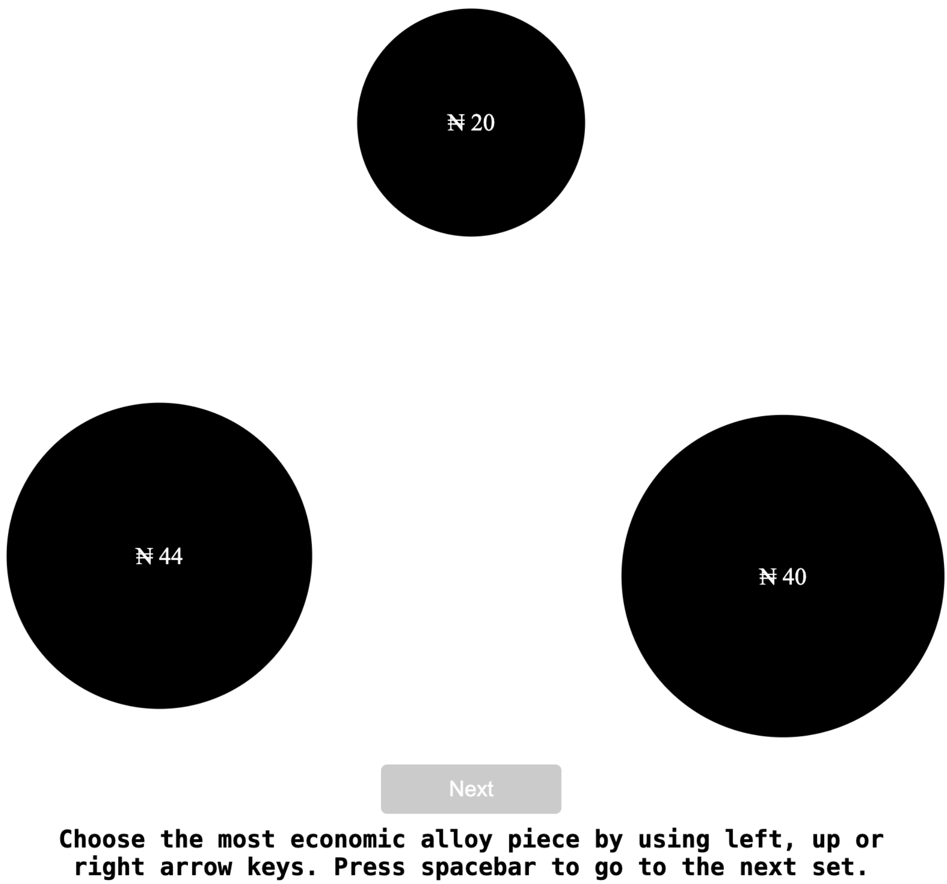

Figure 1. *Example trial in Experiment 1. Each option combined circular area and price, and participants selected the option with the lowest price per unit area..*

2.1.3. Measure

The context effect was quantified using the equal-weights Relative Choice Share of the Target, RST_{EW} (Katsimpokis et al., 2022):

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

where T_X and C_X denote target and competitor selections when the decoy favored option X, and T_Y and C_Y denote the corresponding counts when the decoy favored option Y. A value of .50 indicates no context effect.

2.2. Results

Across participants, mean RST_{EW} was $M = .591$ ($SD = .109$), exceeding the null value of .50. The directional test supported attraction ($t(29) = 4.593, p < .001, d = .839$). A Bayesian one-sample t-test also supported the directional alternative that $RST_{\mathrm{EW}} > .50$ ($\mathrm{BF}_{10} = 273.75$),

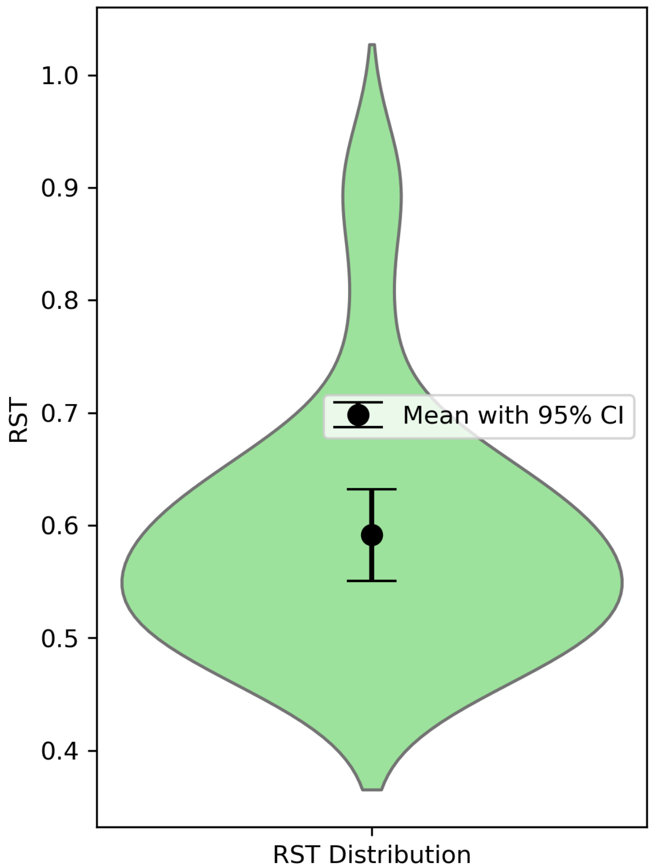

Figure 2. *Distribution of RST_{EW} values in Experiment 1. Values above .50 indicate a relative shift toward the target in the presence of the decoy..*

using a Cauchy prior with $r = .40$. Exploratory response-time analyses did not reveal systematic differences across decoy contexts and are not interpreted further here (Figure 2).

2.3. Discussion

Experiment 1 demonstrated a robust attraction effect in a mixed perceptual–numerical decision task. Participants preferred the target over the competitor when an asymmetrically dominated decoy was introduced, even though the task required integrating circular area with price to compute price per unit area.

The result is consistent with the trade-off-difficulty account, but it does not isolate the source of the effect. The target and competitor were matched on the objective criterion, but that does not reveal whether the observed attraction reflected weak baseline preference between the core options, selective realization of dominance in decoy comparisons, or both. Experiment 2 therefore measured the equivalence geometry underlying this stimulus family.

3. Experiment 2: Adjustment-Derived Equivalence Geometry

Experiment 2 serves as the bridge between the behavioral boundary conditions identified in prior work and the trade-off-difficulty mechanism that can generate them. It is not another test of the attraction effect itself. Instead, it asks whether inter-attribute trade-off difficulty creates the latent equivalence geometry that would appear behaviorally as weak Target–Competitor bias and item-wise Target–Decoy versus Competitor–Decoy asymmetry in pairwise diagnostic tasks.

The ambiguity matters because the present account treats the attraction effect as requiring two separable boundary conditions. The first condition is core-option ambiguity, denoted here as A. The target and competitor must be sufficiently balanced that neither option has a strong pre-decoy advantage. The second condition is item-wise asymmetric dominance, denoted here as S. This condition was part of the original item-level definition of asymmetric dominance proposed by Huber et al. (1982): the decoy should be dominated by the target but not by the competitor. Rath et al. (2025) elaborated this item-wise definition by using a pairwise binary-choice task as a diagnostic for whether a stimulus set actually has this structural property. In that diagnostic, the key question is whether the decoy is rejected in relation to the target more than it is rejected in relation to the competitor. Thus, S is an item-wise property of the stimulus set as revealed by binary choices, not merely an attribute-wise label assigned by the experimenter.

The present experiment asks how such an item-wise structure can be produced in the first place. One possibility is that trade-off difficulty widens the region of subjective equivalence around a reference option. If that widened region includes the shifted target, A is satisfied because the target and competitor are effectively ambiguous. If the same region also includes the decoy, the Competitor–Decoy relation becomes non-diagnostic in the same criterion space, while the Target–Decoy relation remains a direct dominance relation by construction. Under the item-wise definition of asymmetric dominance (Huber et al., 1982), as operationalized by the binary-choice diagnostic in Rath et al. (2025), this geometry is exactly the kind of stimulus structure that should give rise to S: the decoy has a dominance relation to the target that it does not have to the competitor. In that case, trade-off difficulty has not manipulated only one boundary condition. It has simultaneously generated the conditions for both A and S.

The issue is especially important in objective perceptual attraction tasks. In preference tasks, two core options can be balanced subjectively, but there is no single objective criterion that says which option is correct. In perceptual or criterion-based tasks, by contrast, each option has an objective task value. Rectangles can be compared by total area, bars can be compared by total filled length, and circle-price stimuli can be compared by price per unit area. This makes it possible to define a target and competitor that are exactly equivalent, or nearly equivalent, on the objective criterion. However, prior objective attraction studies have not always used perfectly equivalent core options. In designs such as Spektor et al. (2018) and Spektor et al. (2022), core-option discriminability was manipulated using small objective differences, including 3% and 7% offsets. These offsets are not incidental. They can be used to measure pre-decoy bias, prevent participants from recognizing exact equality too easily, and maintain a meaningful objective task structure.

Small offsets create a useful design feature, but they also create a theoretical problem. The present analysis focuses on the 7% target offset because it is the larger, stricter discriminability level from those prior objective tasks. If a participant's equivalence band contains a 7% shift, a smaller 3% shift would also be expected to fall inside it; the reverse need not be true. Whether this 7% offset is psychologically meaningful depends on the participant's equivalence geometry. If the stimulus space is easy to integrate, as in a bar task where the rule is total filled length, then a 7% difference may be discriminable. The target and competitor are not psychologically equivalent, so core-option ambiguity is absent. If the stimulus space is difficult to integrate, as in a circle-price task where a visual area must be related to a symbolic price, then the same 7% difference may fall inside a broad subjective equivalence band. In that case, the target and competitor may be treated as equivalent even though they differ objectively.

The decoy location creates the second part of the problem. In the present stimulus logic, decoys are constructed near the target, approximately 10% from the target. If the participant's equivalence band around the competitor is narrow, this target-referenced decoy location will fall outside the band. The competitor and decoy are then separated in criterion space. But if the participant's equivalence band is broad, the target-referenced 10% decoy may also fall within the competitor-referenced band. In that case, the Competitor–Decoy relation is absorbed into the same equivalence region, whereas the decoy remains attribute-wise dominated by the target by construction. This is the geometry that can produce item-wise asymmetric dominance as diagnosed in binary choice: the decoy has a dominance relation to the target that it does not have to the competitor.

This geometry provides the bridge from trade-off difficulty to the two boundary conditions. If the 7% target offset is inside the competitor-referenced equivalence band, A is satisfied because the target and competitor are ambiguous. If the target-referenced 10% decoy location is also inside that band, then the Competitor–Decoy relation is no longer separated from the competitor in criterion space. Since the decoy is still constructed as dominated by the target, the stimulus set now has the attribute-space arrangement expected to yield S under the Rath et al. (2025) binary-choice diagnostic. In other words, the same broad equivalence region can jointly toggle A and S.

The opposite pattern is also theoretically important. In stimulus spaces that are easy to integrate or nearly commensurate, the equivalence band should be narrower. If both the 7% target offset and the target-referenced 10% decoy location fall outside the band, then neither boundary condition is satisfied in the relevant way. The target and competitor are not ambiguous, so A is absent. The competitor and decoy are also separated in criterion space, so the decoy is not expected to be asymmetrically positioned with respect to the two core options in the item-wise sense. In that case, S is also absent. The expected outcome is a null attraction effect, not a repulsion effect. A negative attraction effect would require a further mechanism that actively favors the competitor or reverses the decoy's influence. A narrow equivalence band by itself predicts that the decoy has little leverage, not that it should push choices away from the target.

This point helps reconcile the present account with the prior bar-stimulus literature. Spektor et al. (2022) reported repulsion-like patterns with bar stimuli, but Rath et al. (2025) reanalyzed those data and concluded that the critical target effect is better characterized as null once the relevant design factors are modeled together. This null pattern is exactly what the present geometry predicts for easy, commensurate stimuli. If participants can compute or approximate total filled length with relatively little subjective dispersion, then target and competitor offsets remain discriminable and the decoy-related comparison structure does not produce the required item-wise asymmetry.

Experiment 2 was designed to test this proposed confound directly. The goal was not to measure another attraction effect. Instead, the goal was to measure the latent equivalence geometry that sits underneath attraction tasks. The experiment used an adjustment paradigm: participants adjusted one attribute of an option until it felt equivalent to a fixed reference under the relevant task rule. For bars, the rule was total filled length. For circle-price stimuli, the rule was value for money, operationalized as price per unit area. Repeated adjustment errors provided an estimate of each participant's equivalence dispersion for each stimulus family. This

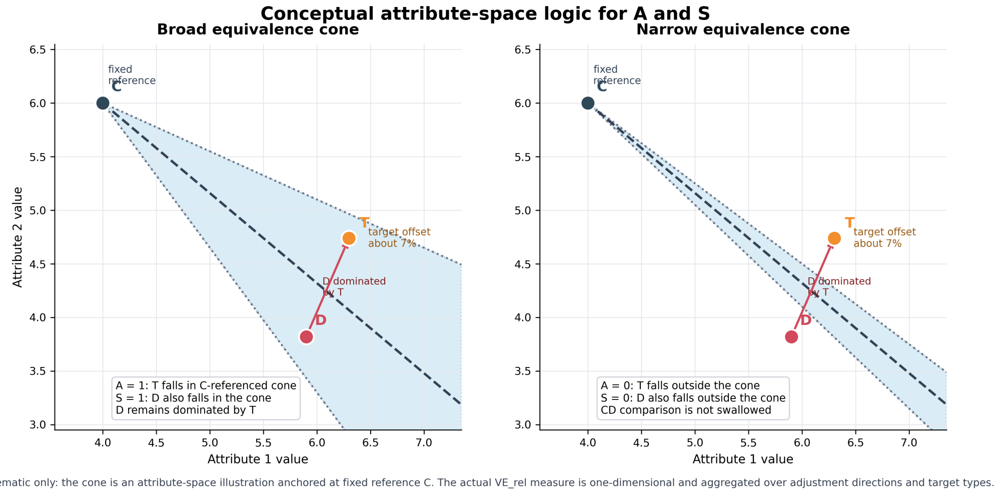

Figure 3. *Conceptual illustration of how adjustment-derived equivalence geometry can jointly toggle A and S. In the adjustment paradigm, the reference option is fixed in both attributes, and the comparison option is adjusted along one attribute at a time until the two options match in criterion value. In a broad competitor-referenced equivalence region, both the target offset and the target-referenced decoy location fall within the region, implying core-option ambiguity and an ambiguous Competitor–Decoy comparison while Target–Decoy dominance remains comparatively easy. In a narrow region, both offsets fall outside the region. The figure is schematic: the empirical analysis projects these adjustment responses onto participant-level VE_{rel} in criterion-error space..*

participant-specific dispersion was then used to define an equivalence band. The critical question was whether the preregistered 7% target offset and target-referenced 10% decoy location fell inside that band (Figure 3).

3.1. Method

3.1.1. Participants and Sampling Plan

Participants were recruited from the institutional participant pool and provided informed consent. Data collection followed a preregistered sequential Bayesian sampling plan with a minimum sample size of 20 and maximum sample size of 70. Sampling was to stop if evidence for the alternative reached $\mathrm{BF}_{10} \geq 10$, if evidence for the null reached $\mathrm{BF}_{01} \geq 10$, or if $N_{\max} = 70$ was reached. Twenty-two participants completed the study. Two were excluded according to the preregistered catch-trial criterion, leaving $N = 20$.

3.1.2. Design and Procedure

The experiment used a fully within-participant design. Each participant completed 80 bar adjustment trials and 80 circle-price adjustment trials, along with 20 catch trials. Block order was counterbalanced. On each trial, a fixed reference option appeared on the left and an adjustable option appeared on the right. Participants used a slider to adjust one attribute of the adjustable option until it matched the reference in overall value under the displayed rule, then confirmed their response (Figure 4).

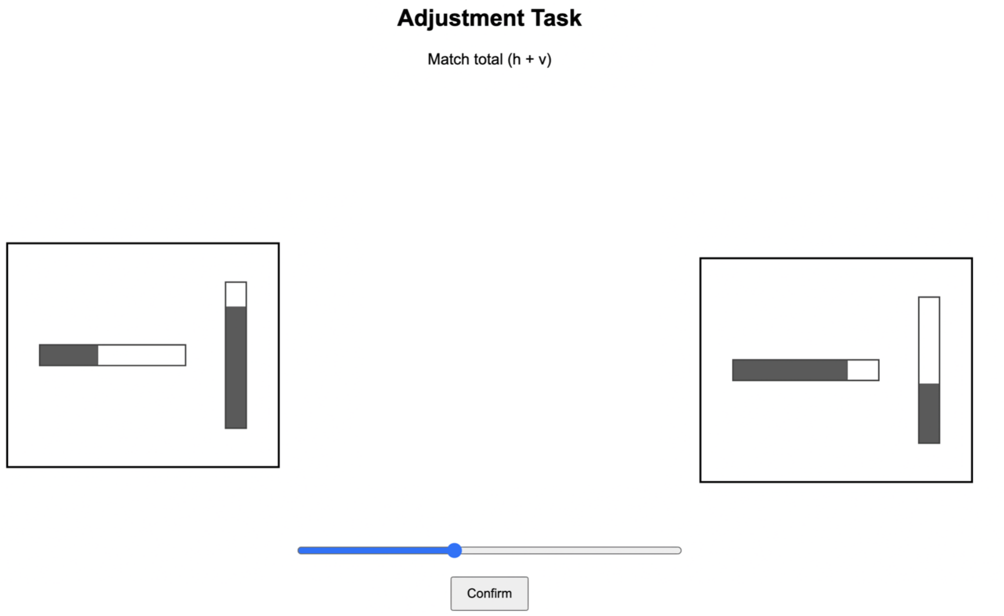

Figure 4. *Example adjustment trials from Experiment 2. Participants adjusted one attribute of the option on the right until it matched the fixed reference under the task rule. For bars, equivalence was defined by total filled length; for circle-price stimuli, equivalence was defined by price per unit area..*

3.1.3. Equivalence Dispersion

Each adjustment trial had a true match value. Let f_A denote the task-relevant criterion value of the fixed reference and f_B the criterion value of the adjusted option. For bars, f was total filled length; for circle-price stimuli, f was price per unit area. The relative criterion error was

$$e_{\text{rel}} = \frac{f_B - f_A}{f_A}. \tag{2}$$

For each participant and stimulus family, equivalence dispersion was defined as

$$VE_{\text{rel}} = SD(e_{\text{rel}}). \tag{3}$$

The preregistered equivalence band was $k \times VE_{\text{rel}}$, with $k = 1.5$. An offset was classified as included if its magnitude did not exceed this band. Figure 5 illustrates the construction.

3.1.4. Structural Classification

The structural analysis evaluated two predefined points: a 7% target offset and a target-referenced 10% decoy location. If the 7% target offset fell inside a participant's equivalence band, the target and competitor were treated as psychologically equivalent at that offset. If the target-referenced decoy location also fell inside the same competitor-referenced band, the Competitor–Decoy relation was treated as absorbed by the equivalence region. Participants were classified as showing joint structural inclusion when both offsets fell inside the band.

H1 compared equivalence dispersion between stimulus families with a paired t-test and a Bayesian paired t-test. H2, the primary confirmatory analysis, compared paired joint-inclusion classifications using McNemar's test and a preregistered Bayesian paired binary analysis based

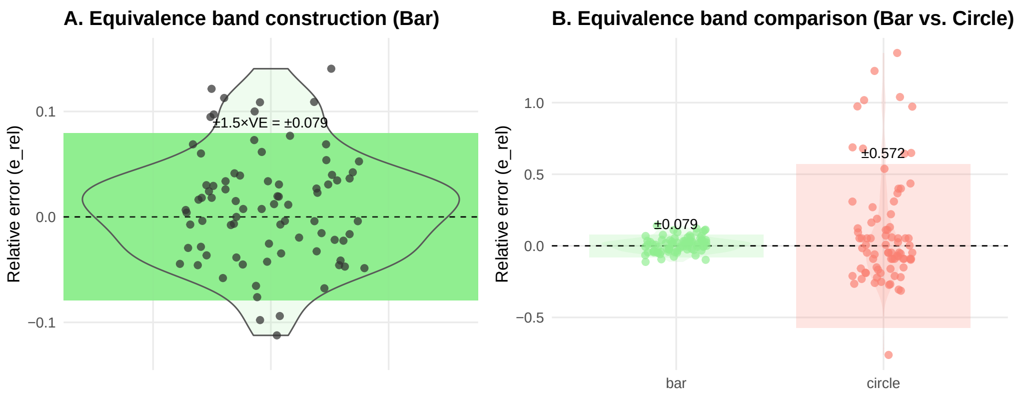

Figure 5. *Illustration of equivalence band estimation in Experiment 2. Each point represents a single adjustment response expressed as relative error* (e_{rel})*. The shaded region denotes the participant-specific equivalence band,* $\pm k \cdot VE_{\text{rel}}$*, with* $k = 1.5$*..*

on discordant counts. Robustness checks varied k from 1.0 to 2.0 and examined continuous standardized distances from the equivalence boundary.

3.2. Results

3.2.1. H1: Equivalence Dispersion

Circle-price stimuli exhibited greater equivalence dispersion than bar stimuli ($t(19) = 3.504, p = .0024, d_z = .783$). The corresponding Bayesian paired t-test, using a Cauchy prior with $r = .707$, yielded strong evidence that dispersion was greater for circle-price stimuli than for bar stimuli ($\text{BF}_{10} = 16.972$).

3.2.2. H2: Joint Inclusion of Target and Decoy Offsets

Circle-price stimuli produced joint inclusion in 100% of participants, whereas bar stimuli produced joint inclusion in 25% of participants. The paired risk difference was .750. A McNemar test supported the paired classification difference ($\chi^2(1) = 13.067, p = .0003$). The preregistered Bayesian paired binary analysis yielded extremely strong evidence that circle-price stimuli produced greater joint inclusion than bar stimuli ($\text{BF}_{10} = 2048.000$). The discordant pattern was 15 circle-only classifications and 0 bar-only classifications.

The robustness analyses supported the same conclusion. Continuous standardized distances showed that decoy locations lay closer to the equivalence boundary for circle-price than bar stimuli. Varying the equivalence-band multiplier from 1.0 to 2.0 preserved the positive circle-minus-bar difference; the proportion difference ranged from .542 to .792, and BF_{10} ranged from 585.14 to 26214.40. Full robustness plots are reported in the supplementary materials.

3.3. Discussion

Experiment 2 showed that the same stimulus family that produced attraction in Experiment 1 also produced broad equivalence regions. For every retained participant, circle-price stimuli yielded equivalence bands broad enough to include both the target offset and target-referenced decoy location. Bar stimuli produced the same joint inclusion for only one quarter of participants.

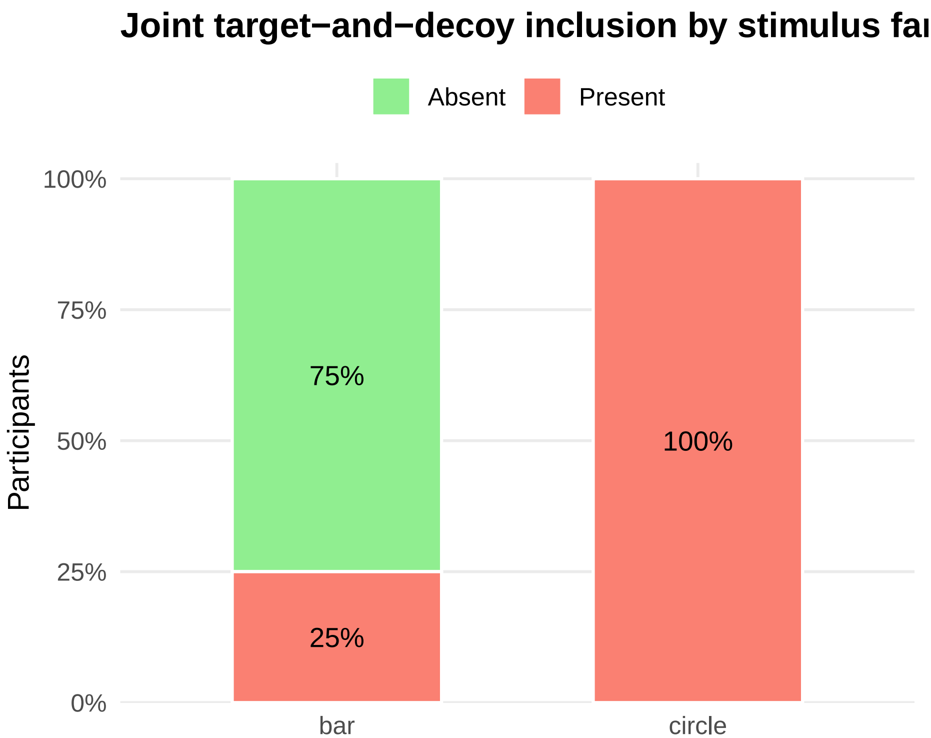

Figure 6. *Gate configuration for joint target-and-decoy inclusion in Experiment 2. Bars show the proportion of participants whose equivalence band jointly included both the 7% target offset and the 10% target-referenced decoy location..*

This finding clarifies the trade-off-difficulty account. The relevant link is not difficulty in the abstract, but the way difficulty reshapes the comparison geometry of the choice set. A broad equivalence region can make the target and competitor ambiguous while also making the Competitor–Decoy relation non-diagnostic, leaving the Target–Decoy dominance relation comparatively clear. Thus, trade-off difficulty can move both A and S together.

4. General Discussion

Across two experiments, we found that a mixed perceptual–numerical stimulus family can produce a robust attraction effect and that the same family also broadens subjective equivalence geometry in a way that bundles two conditions for attraction. Experiment 1 established attraction in a circle-area/price choice task. Experiment 2 showed that circle-price stimuli produced broader equivalence dispersion than bar stimuli and, more importantly, placed the target offset and target-referenced decoy location inside participants' equivalence bands much more often than bars.

The contribution is diagnostic. The results explain why difficult trade-offs may be favorable to attraction without showing that trade-off difficulty is a single causal switch. When a stimulus family is difficult to integrate, it may simultaneously create room for the decoy to matter and structure decoy-related comparisons in a way that favors the target. Conversely,

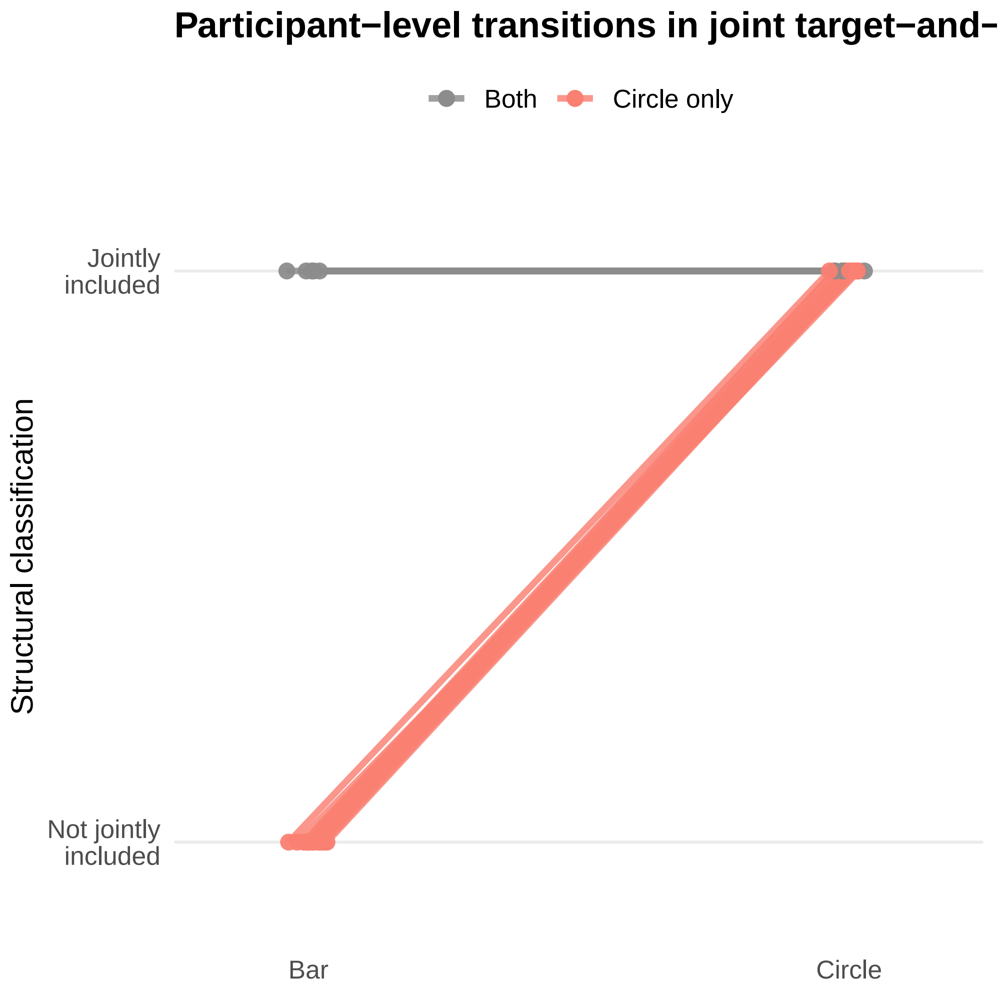

Figure 7. *Paired participant-level transition in joint target-and-decoy inclusion. Each line represents one participant's classification under the bar and circle-price stimulus families..*

near-commensurate stimuli may remove both ingredients at once. Positive-versus-null contrasts between these stimulus families are therefore descriptively meaningful but not causally decisive.

This distinction helps explain why the attraction effect can be robust in some settings and elusive in others. The effect depends not only on nominal asymmetric dominance, but on whether the relevant comparison geometry is psychologically realized. The present experiments show that subjective equivalence regions provide one route by which trade-off difficulty can create that geometry. A separate off-diagonal design is required to test whether core-option ambiguity alone is sufficient or whether decoy-related comparative asymmetry is necessary once ambiguity is in place. The current paper therefore sets up that causal question while keeping its own claim deliberately narrower: trade-off difficulty can bundle the ingredients that make attraction possible.

Acknowledgments. Acknowledgments are masked for review.

Funding Statement. Funding information is masked for review. If no specific funding supported this work, replace this paragraph with: This research received no specific grant from any funding agency, commercial, or not-for-profit sectors.

Competing Interests. The author(s) declare none.

Data Availability Statement. All data, analysis code, experimental materials, figure-generation scripts, and analysis outputs needed to reproduce the reported results will be made available on the Open Science Framework at the time of submission. The final submission should replace this sentence with the public repository URL.

Ethical Standards. All procedures were approved by the relevant institutional ethics committee, and all participants provided informed consent before participation.

Supplementary Materials. Supplementary materials include equivalence-band sensitivity analyses, participant-level equivalence dispersion, continuous distance analyses, and sequential Bayes factor trajectories for Experiment 2.

Author Contributions. Author contributions are masked for review.

References

Banerjee, P., Hayes, W. M., Chatterjee, P., Masters, T., Mishra, S., & Wedell, D. H. (2024). Factors that promote the repulsion effect in preferential choice. *Judgment and Decision Making*, *19*, e11.

Chang, R. (2013). Incommensurability (and incomparability).

Hayes, W. M., Holmes, W. R., & Trueblood, J. S. (2025). Attribute commensurability and context effects in preferential choice. *Psychonomic Bulletin & Review*, *32*(2), 737–748.

Hedgcock, W., & Rao, A. R. (2009). Trade-off aversion as an explanation for the attraction effect: A functional magnetic resonance imaging study. *Journal of Marketing Research*, *46*(1), 1–13.

Huber, J., Payne, J. W., & Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*.

Huber, J., Payne, J. W., & Puto, C. P. (2014). Let's be honest about the attraction effect. *Journal of Marketing Research*, *51*(4), 520–525.

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

Luce, M. F., Bettman, J. R., & Payne, J. W. (2001). Emotional decisions: Tradeoff difficulty and coping in consumer choice. *Monographs of the journal of consumer research*, (1), 1–209.

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

Pettibone, J. C. (2012). Testing the effect of time pressure on asymmetric dominance and compromise decoys in choice. *Judgment and Decision making*, *7*(4), 513–521.

Rath, T., Srinivasan, N., & Srivastava, N. (2025). The attraction effect in perceptual decision-making: A case of dominance asymmetry. *Frontiers in Psychology*, *16*, 1661748.

Rath, T., Srivastava, N., & Srinivasan, N. (2026). *Unmasking the flaws of triplet-triplet attraction effect measures: Via mathematical analysis and agent-level simulations* [Manuscript in preparation].

Simonson, I., & Tversky, A. (1992). Choice in context: Tradeoff contrast and extremeness aversion. *Journal of marketing research*, *29*(3), 281–295.

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., Kellen, D., & Hotaling, J. M. (2018). When the good looks bad: An experimental exploration of the repulsion effect [Publisher: Sage Publications Sage CA: Los Angeles, CA]. *Psychological science*, *29*(8), 1309–1320.

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., & Busemeyer, J. R. (2013). Not just for consumers: Context effects are fundamental to decision making. *Psychological science*, *24*(6), 901–908.

Tversky, A., & Simonson, I. (1993). Context-dependent preferences. *Management science*, *39*(10), 1179–1189.

Walasek, L., & Brown, G. D. (2024). Incomparability and incommensurability in choice: No common currency of value? *Perspectives on Psychological Science*, *19*(6), 1011–1030.

Supplementary Materials for "When Difficult Trade-Offs Make Decoys Work: Subjective Equivalence Geometry in the Attraction Effect"

Author names masked for review

S1: Sensitivity Analysis of the Equivalence Band

The structural ambiguity analysis in Experiment 2 defined joint inclusion of the target and decoy offsets relative to each participant's empirical equivalence dispersion, VE_{rel}, using bands of the form

$$|\Delta| \leq k \cdot VE_{\text{rel}}.$$

Although the preregistered analysis used $k = 1.5$, we assessed robustness by recomputing gate classifications across $k \in [1.0, 2.0]$ in steps of 0.05, holding the preregistered offsets $\Delta_T = 0.07$ and $\Delta_D = -0.10$. This analysis tests whether structural inclusion of both offsets depends on the equivalence threshold.

For each k and stimulus family, we computed whether both offsets lay inside the band. We report the proportion of participants satisfying the joint gate under circle-price stimuli minus the corresponding proportion under bar stimuli, and a paired-binary Bayes factor from discordant classifications.

Across the tested grid, the proportion difference remained positive for all k. The proportion difference ranged from .542 to .792, and BF_{10} ranged from 585.14 to 26214.40.

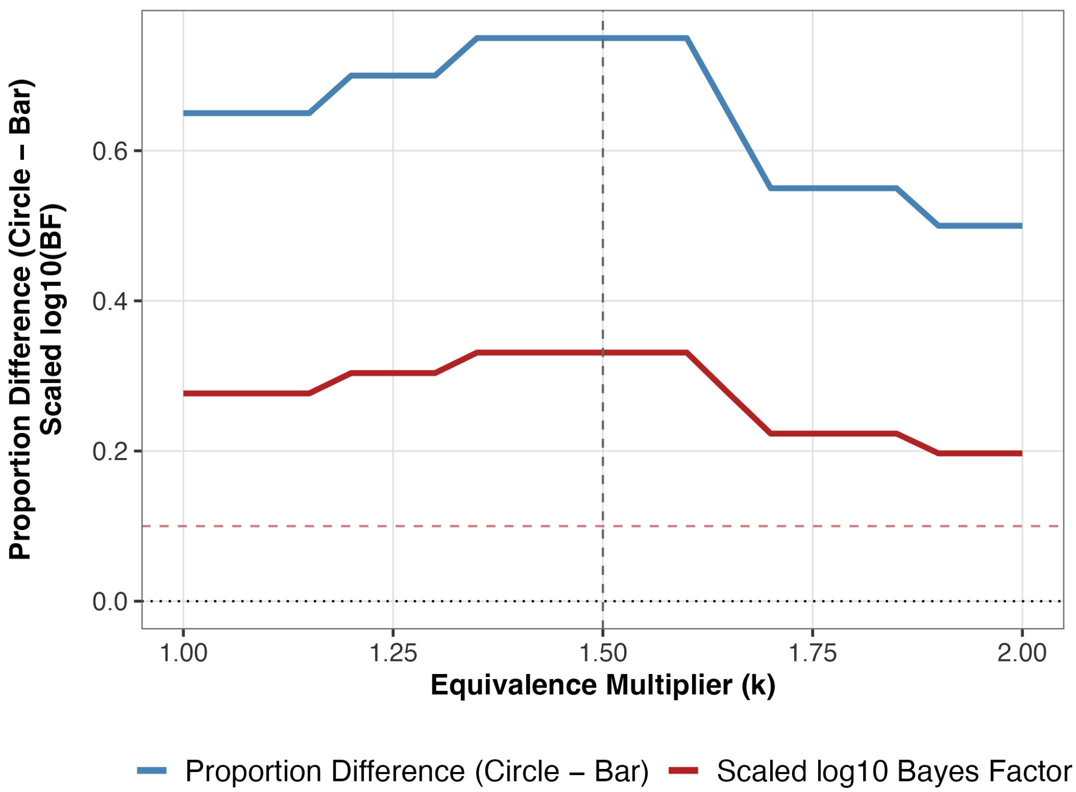

Figure 1: **Experiment 2 H2 sensitivity to the equivalence multiplier k.** The solid blue line shows the circle-minus-bar proportion difference in joint inclusion. The solid firebrick line shows scaled $\log_{10}(\text{BF}_{10})/10$. The dashed grey vertical line marks the preregistered $k = 1.5$. Across the grid, the proportion difference remained positive and Bayes factors remained strong.

S2: Equivalence Dispersion by Stimulus Family

The main text reports the confirmatory paired comparison of equivalence dispersion. Figure 2 shows participant-level equivalence dispersion for bar and circle-price stimuli.

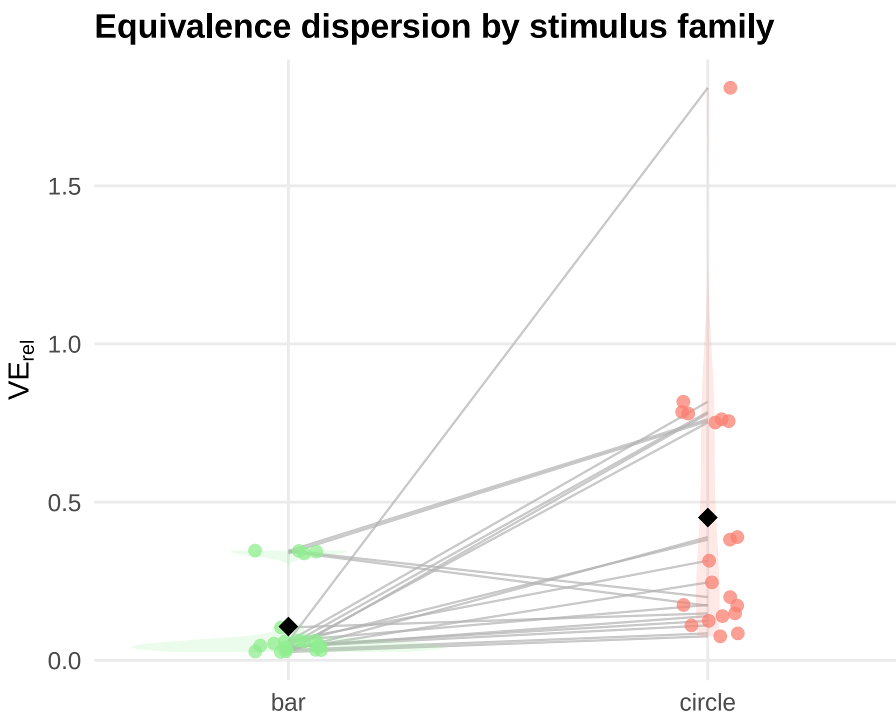

Figure 2: *Equivalence dispersion in Experiment 2 by stimulus family.* Each connected pair represents one participant's VE_{rel} for bar and circle-price stimuli. Larger values indicate broader participant-level equivalence regions in criterion-error space.

S3: Continuous Distance Analysis

To assess whether the structural effect depends on the binary inclusion criterion, we conducted a complementary analysis using a continuous measure of proximity to the equivalence boundary. For each participant and stimulus type, the target and decoy offsets were expressed in units of the participant's equivalence dispersion:

$$d = \frac{|\Delta|}{VE_{\text{rel}}}.$$

Smaller values indicate that an offset lies closer to the participant's equivalence region. Values below the chosen cutoff indicate inclusion.

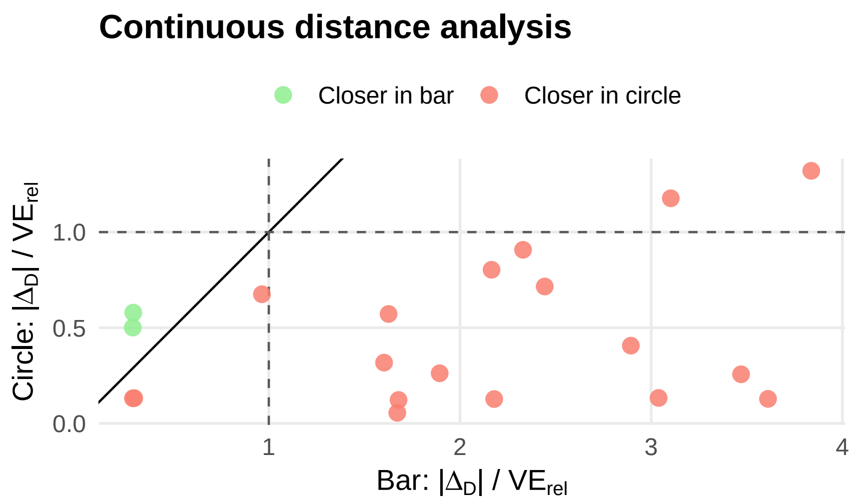

Figure 3: *Continuous robustness analysis for Experiment 2.* Each point represents a participant's standardized decoy distance, $|\Delta_D|/VE_{\text{rel}}$, for bar stimuli (x-axis) and circle-price stimuli (y-axis). The diagonal indicates equality. Points below the diagonal indicate that the decoy lies closer to the equivalence boundary under circle-price stimuli. Dashed lines at $d = 1$ denote the equivalence boundary.

The distribution showed the same qualitative pattern as the binary gate analysis: decoy offsets were closer to the equivalence boundary for circle-price stimuli than for bar stimuli.

S4: Sequential Bayes Factor Trajectories

To assess evidence accumulation over sample size, we computed sequential Bayes factors for both preregistered hypotheses using prefix samples, incrementally adding participants from

$n = 3$ to the final sample. At each step, Bayes factors were recomputed using all participants up to that sample size.

For H1, Bayes factors were computed from the paired-difference model comparing VE_{circle} and VE_{bar}, using the preregistered Cauchy prior on standardized effect size. For H2, Bayes factors were computed using a paired binary beta-binomial model applied to discordant gate classifications.

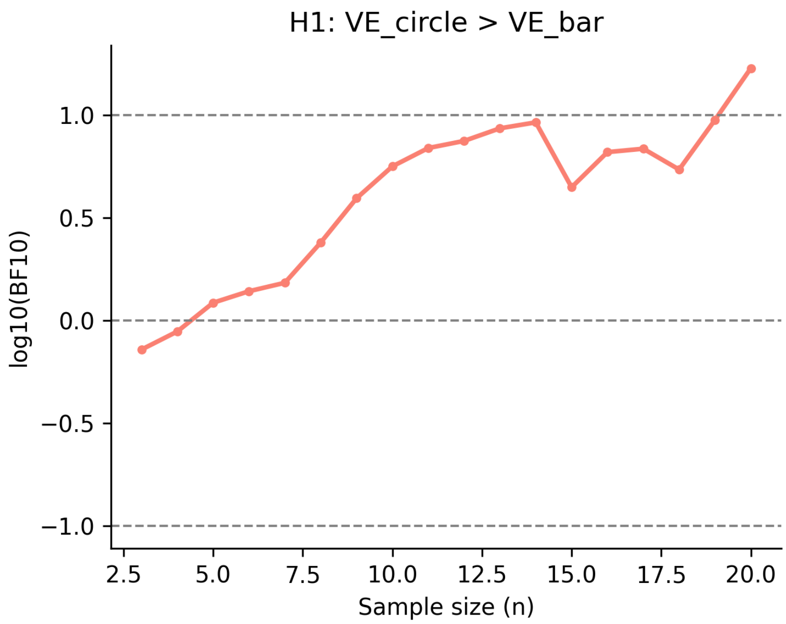

Figure 4: *Sequential Bayes factor trajectory for H1 in Experiment 2.* Bayes factors for the hypothesis $VE_{\text{circle}} > VE_{\text{bar}}$ are shown as a function of cumulative sample size.

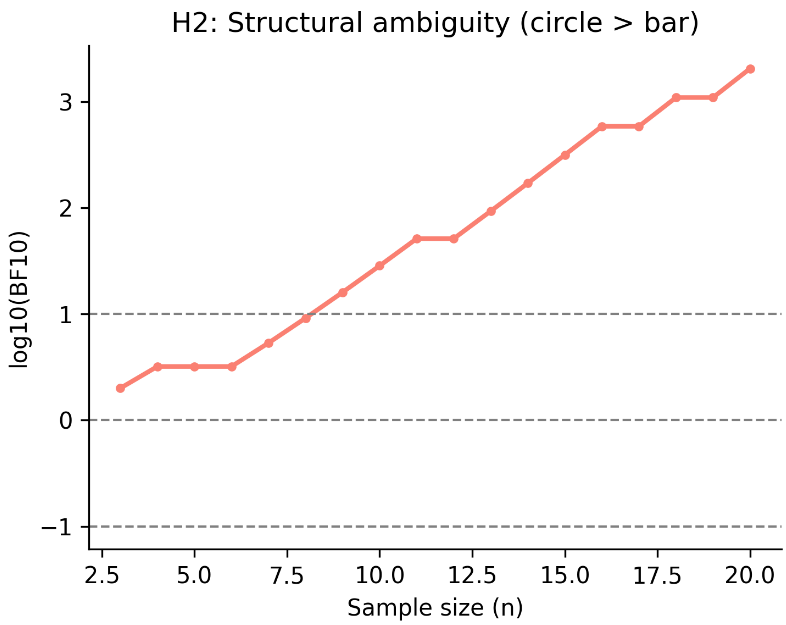

Figure 5: *Sequential Bayes factor trajectory for H2 in Experiment 2.* Bayes factors for greater joint structural inclusion under circle-price relative to bar stimuli are shown as a function of cumulative sample size.