Cognition in Preferences and Choice

Maria Betto

February · 2026

Abstract

Observed choices often become less selective when decision-makers are distracted, rushed and/or fatigued, even when underlying tastes are stable. We develop a theory of cognition-indexed behavior consisting of a family of choice correspondences, and propose a cognition-dependent representation (CDR) with a (i) fixed benchmark utility, representing the decision-maker's stable tastes, and (ii) cognition-dependent thresholds that determine how much suboptimality they are willing to tolerate. We characterize this model behaviorally via menu-wise inclusion (WOCI) and a cross-cognition axiom (WARPD) that generalizes the standard Weak Axiom; jointly, they are necessary and sufficient for a CDR. We further show equivalence to rationalizability by a homogeneous nested family of semiorders satisfying Roberts-type axioms, and to a fuzzy representation in which a single fuzzy difficulty relation and an order-preserving calibration generate the observed correspondences as cut choices.

1 Introduction

Observed choices often vary systematically with the decision environment. The same person may make a sharp selection when stakes are salient and time is ample, yet treat several options as interchangeable when distracted, rushed, fatigued, or facing noisy information that blurs fine comparisons. A natural interpretation is not that tastes shift across environments, but that the agent's ability (or willingness) to act on small differences changes. When discrimination is coarse, suboptimal alternatives may still be "good enough" and remain admissible; when discrimination is fine, admissible sets contract toward the exact maximizers of a stable underlying ranking.

This paper develops a revealed-preference theory of such cognition-indexed choice behavior. We start from a family of choice correspondences $\mathcal{C} = \{C_\lambda : \lambda \in \Lambda\}$, where each C_λ assigns to every menu S a nonempty subset $C_\lambda(S) \subseteq S$ of *admissible* options under cognition level λ. The index λ summarizes reduced-form features of the environment that affect discrimination—attention, time pressure, incentives, expertise, noise, and so on—without presuming that cognition regimes are ordered a priori. Our question is: when does an observed family $\mathcal{C}$ arise from a single stable preference ranking together with cognition-dependent limits on discrimination?

To that end, we introduce the following representation. We say that $\mathcal{C}$ admits a *cognition-dependent representation* (CDR) (Definition 1) if there exist a utility index $u : X \to \mathbb{R}$ encoding a stable benchmark ranking, a tolerance profile $\varepsilon(\lambda, x) \geq 0$, and a weak order on cognition indices with a maximal element $\overline{\lambda}$, such that for every λ and menu S,

$$C_\lambda(S) = \{x \in S : \text{ for all } y \in S,\ u(y) - u(x) \leq \varepsilon(\lambda, x)\}.$$

Tolerances shrink with cognition, the endpoints $u(x)+\varepsilon(\lambda, x)$ respect the benchmark ranking, and at maximal cognition tolerances vanish: $\varepsilon(\overline{\lambda}, x) = 0$ for all x. In this representation, cognition is neither a taste shifter nor a welfare primitive; its role is to determine which benchmark differences are behaviorally actionable.

The CDR is closely related to the classical semiorder (or just-noticeable-difference) paradigm. For a single static environment, semiorders provide a standard JND (just-noticeable difference) account of coarse discrimination (Luce, 1956, Fishburn, 1970b). A choice correspondence C is semiorder-rationalizable if there exist $v : X \to \mathbb{R}$ and $\theta \geq 0$ such that

$$C(S) = \{x \in S : \forall y \in S, v(y) - v(x) \leq \theta\}.$$

A constant gap θ is, however, not well suited for comparing discrimination across environments when the benchmark ranking is treated as ordinal: $\arg\max$ behavior is invariant to monotone relabelings of utility, while a constant additive gap is meaningful only up to affine rescalings. As a result, modeling cognition as changes in a single scalar θ implicitly fixes a utility metric and can impose spurious cardinal restrictions; Section 2 provides an explicit example.

The CDR addresses this concern by construction. It is fully ordinal in the sense that the benchmark index u is identified only up to strictly increasing transformations of u. Any monotone relabeling of u can be offset by a corresponding transformation of the tolerances ε, leaving $\{C_\lambda : \lambda \in \Lambda\}$ unchanged. In this way, the representation commits to a stable *ranking* while remaining agnostic about cardinal magnitudes: cognition affects admissibility through thresholds that may shrink non-uniformly across alternatives, rather than through a single scalar measured in fixed utility units.

Our first contribution is to provide an axiomatic characterization of CDR in terms of choice behavior. Axiom WOCI recovers a weak order on Λ from menu-wise inclusion: $\lambda' \geq \lambda$ if and only if $C_{\lambda'}(S) \subseteq C_\lambda(S)$ for all menus S. We then impose a cross-cognition revealed-preference axiom, WARPD. It weakens standard WARP by anchoring comparisons at alternatives that remain admissible under maximal selectivity. Informally, if x survives maximal cognition in S, then at any lower cognition level, whenever some y is admissible from a menu T that also contains x, two consistency requirements hold: (i) x must be admissible from T, and (ii) y must be admissible from S.

Theorem 1 shows that these two axioms are exactly what is needed: $\mathcal{C}$ satisfies WOCI and WARPD if and only if it admits a CDR. Thus the cognition-indexed admissibility sets are rationalizable by a single benchmark ranking together with cognition-dependent tolerances.

Next, we provide a preference-side characterization. For each cognition level λ, define a strict relation P_λ from binary menus, where $xP_\lambda y$ means that at cognition λ the agent recognizes x as strictly better than y. Theorem 2 shows that CDR is equivalent to rationalizability by a homogeneous family of semiorders satisfying certain properties (GA, N and UW). These conditions were taken from Roberts (1971), who were interested in studying stochastic-data problems and their probabilistic consistency.

We then show an equivalent fuzzy representation. The entire system can be encoded by a single *fuzzy-rational* relation $R : X \times X \to V$, where $R(x, y)$ measures the difficulty of recognizing the strict benchmark statement "$y \succ x$." Cognition levels correspond to resolution thresholds via an order-preserving calibration $\nu : \Lambda \to V$ normalized by $\nu(\overline{\lambda}) = \bar{v}$. Equivalently, CDR is characterized by *fuzzy rationalizability*: there exist a calibration ν and a fuzzy-rational R (fuzzy complete and fuzzy transitive) such that each observed C_λ coincides with the $\nu(\lambda)$-cut choice correspondence induced by R. In this sense, our model is a fully ordinal fuzzy extension of standard rational choice.

This equivalence also clarifies a link between two classic responses to the Sorites paradox: semiorders allow nontransitive local indistinguishability, while fuzzy sets allow graded truth. Under our calibration, these coincide as the same cut structure (Section 5.5).

Finally, we study an extension. Section 7 treats incomplete observation of cognition regimes and constructs a revealed lower bound on the (possibly unobserved) benchmark strict order.

The paper is organized as follows. Section 2 introduces CDR and motivates the separation between ordinal benchmark information and discrimination parameters. Sections 2–4 provide the choice-side and preference-side axioms and prove the main equivalence results (Theorems 1 and 2). Section 5 develops fuzzy rationalizability. Section 7 presents the extension.

2 Setup

Let X be a finite set of alternatives and let $\mathcal{S}$ denote the set of all nonempty subsets of X (the set of all menus). We observe a family of nonempty-valued choice correspondences $\mathcal{C} := \{C_\lambda : \lambda \in \Lambda\}$, where each $C_\lambda : \mathcal{S} \to \mathcal{S}$ satisfies $C_\lambda(S) \subseteq S$ for all $S \in \mathcal{S}$.

We refer to each $\lambda \in \Lambda$ as a *cognition index*. We also say that the choice correspondence $C_\lambda(S)$ is the set of *admissible* options from menu S at cognition level λ.

At this stage, Λ is simply a set of "labels" λ: we impose no a priori order or structure on this set.

In applications, λ can be understood as a reduced-form descriptor of forces that make fine discrimination easier or harder (focus, time pressure, expertise, noise, fatigue, incentives, and so on). The analyst may observe treatments or covariates that classify observations into different regimes indexed by λ.

Our construction will rely on the idea that the decision-maker has a stable underlying ranking of alternatives (encoded by an utility index u, which will be introduced below), but that their ability or willingness to act on small differences of this utility varies with cognition. Low cognition means that small utility advantages are behaviorally negligible, so alternatives that are not truly optimal may still be admissible because they are close enough to the menu-best option. High cognition means finer discrimination and therefore greater selectivity: admissible sets shrink toward the exact maximizers of the underlying ranking.

Importantly, cognition is not a primitive taste shifter and is not meant to proxy welfare or experienced utility. The framework is agnostic about its foundations: a given level of cognition may reflect perceptual limits, scarcity of attention, limited domain knowledge, or low deliberation effort.[1]

Note that the index λ need not be directly observed, as this might be impossible ('how hard is the decision-maker thinking?' is hard to assess). It is however required that the analyst be able to classify observations into a set of "cognition regimes" and treat these regimes as distinct from each other in terms of discriminability. This may come from experimental design (for example, explicit treatments that increase time, attention, or incentives), or from an empirical proxy that plausibly affects discriminability (for example, noise level, measured distraction, or expertise categories). This cognition regimes can be classified into rankable "cognition levels" only if certain conditions hold.

Finally, note that this interpretation is consistent with a long tradition of threshold-based models of imperfect discrimination: preferences are not assumed to change, but the decision-maker behaves as if differences below a tolerance band are not reliably actionable.

2.1 Cognition-dependent utility representation

Recall the standard utility-maximization model. A choice correspondence C is (weak-order) *rationalizable* if there exists a utility function $u : X \to \mathbb{R}$ such that, for every $S \in \mathcal{S}$,

$$C(S) = \operatorname{argmax}_{x \in S} u(x) = \{x \in S : \text{ for all } y \in S,\ u(y) - u(x) \leq 0\}. \tag{1}$$

This is the perfect-discrimination benchmark: whenever y yields even slightly higher utility than x (in whatever utility scale), the decision-maker can reliably tell and therefore rules x out.

A classical weakening replaces exact optimality by a just-noticeable-difference (JND) threshold. We say that C is *semiorder–rationalizable* if there exist $v : X \to \mathbb{R}$ and $\theta \geq 0$ such

[1] If cognitive effort is itself chosen, the model can be read conditionally: holding fixed the realized cognitive state or environment, observed choices satisfy the threshold structure below, without taking a stand on the costs or incentives that generated that state.

that, for every $S \in \mathcal{S}$,

$$C(S) = \{x \in S : \text{ for all } y \in S,\ v(y) - v(x) \leq \theta\}. \tag{2}$$

Thus x is admissible whenever it lies within θ units of the menu optimum. The case $\theta = 0$ reduces to (1).

Expression (2) appears as a quantitative measure of "how close" the decision-maker is to exact maximization. In most choice environments, however, the numerical magnitude of θ has no standalone meaning: if (v, θ) satisfies (2), then (v', θ') where $v'(x) = av(x) + b$ and $\theta' = a\theta$ also satisfies (2), for any $a > 0$. Hence θ can always be normalized (for example to $\theta = 1$) by rescaling the utility axis. This is innocuous when studying a single correspondence in isolation, but becomes problematic once one wants to compare *different* cognition levels to a common benchmark utility u.

Furthermore, the benchmark utility in (1) identifies only an ordinal ranking (any strictly increasing transformation preserves argmax), while a constant additive gap implicitly selects a particular cardinalization.

The key point is that semiorder choice is invariant to positive affine rescalings of v (which preserve a constant gap), but not invariant to general monotone relabelings (which preserve only the benchmark ranking). Consequently, if the benchmark utility is understood as an ordinal preference ranking, constant-gap semiorder representations are not comparable across cognition levels without additional cardinal structure. The example below evidences the issue.

Example 1. An agent repeatedly hears two tones that differ only in physical intensity (sound power per unit area), measured in mW/m^2. The agent must declare which tone is louder. Let $X = \{a, b, c, d, e\}$ index five possible intensities, and assume the agent's stable "preference" is simply more intense is louder.

At acuity level λ, the agent can reliably perceive that x is louder than y only if the absolute intensity gap exceeds a constant just-noticeable-difference (JND) threshold $\theta > 0$:

$$v(x) - v(y) \geq \theta,$$

where $v(x)$ is the physical intensity of stimulus x. Take for example

$$v(a) = 1.00, \quad v(b) = 1.04, \quad v(c) = 1.80, \quad v(d) = 1.84, \quad v(e) = 1.86, \qquad \theta = 0.08.$$

Then b does not strictly dominate a because $v(b) - v(a) = 0.04 < 0.08$, and d does not strictly dominate c because $v(d) - v(c) = 0.04 < 0.08$. Also e does not strictly dominate c, since $v(e) - v(c) = 0.06 < 0.08$.

At a higher acuity level, say λ', suppose the JND is proportional to the baseline intensity (a Weber law, Weber, 1834): the agent perceives x louder than y if and only if the proportional increase exceeds $\theta' > 0$,

$$\frac{v(x) - v(y)}{v(y)} > \theta' \quad \Longleftrightarrow \quad \frac{v(x)}{v(y)} > 1 + \theta'.$$

Fix $\theta' = 0.03$ and restrict attention to the reasonable intensity range $v(\cdot) \in [1, 2]$. Then the required absolute gap under λ' is $\theta' v(y) \leq 0.03 \cdot 2 = 0.06 < \theta$, so λ' is strictly more discriminating than λ throughout this range.

With the same five stimuli, we get:

$$\frac{1.04}{1.00} = 1.04 > 1.03, \qquad \frac{1.84}{1.80} \approx 1.022 < 1.03,$$

so under λ' we have b strictly dominating a while d still does not strictly dominate c. Moreover e strictly dominates c, since

$$\frac{1.86}{1.80} \approx 1.033 > 1.03.$$

Thus λ' distinguishes (b, a) and (e, c) even though λ does not, while it still fails to distinguish (d, c).

Define $u := \ln \circ v$ and set

$$\tau := \ln\big(1 + \theta'\big) = \ln(1.03).$$

Then for any x, y,

$$u(x) - u(y) \geq \tau \Longleftrightarrow \ln v(x) - \ln v(y) \geq \ln(1.03) \Longleftrightarrow \frac{v(x)}{v(y)} \geq 1.03,$$

so (u, τ) is a constant-threshold semiorder representation of the proportional discrimination rule. (Equivalently, this is a constant JND on a decibel-type scale, since dB $\propto \log v$.)

Suppose, toward a contradiction, that choices at λ' could be represented on the original intensity scale by some constant absolute threshold $\tilde{\theta} > 0$:

$$\frac{v(x) - v(y)}{v(y)} > \theta' \Longleftrightarrow v(x) - v(y) \geq \tilde{\theta}.$$

Because λ' distinguishes (b, a), we must have

$$\tilde{\theta} \leq v(b) - v(a) = 0.04.$$

Because λ' does *not* distinguish (d, c), we must have

$$\tilde{\theta} > v(d) - v(c) = 0.04.$$

This is impossible. The proportional rule distinguishes relative gaps, so it can treat two equal absolute gaps differently depending on the baseline level, and therefore it cannot be represented by a constant absolute-gap threshold on the raw intensity scale. $\lozenge$

A convenient way to separate what is behaviorally meaningful from what is mere scaling is to rewrite the threshold model so that the benchmark utility remains purely ordinal and all gap information is moved into the threshold.

To do so, take any semiorder representation (v, θ) of C and apply a strictly increasing relabeling $f : \mathbb{R} \to \mathbb{R}$. Define $u(x) \coloneqq f(v(x))$. Then u represents the same underlying preference ranking as v, but the threshold becomes alternative-specific:

$$v(y) - v(x) \leq \theta \iff u(y) - u(x) \leq \underbrace{f(v(x) + \theta) - f(v(x))}_{=:\theta(x)}.$$

This yields an equivalent formulation: there exist a utility mapping $u : X \to \mathbb{R}$ and a threshold function $\theta : X \to \mathbb{R}_+$ such that, for every $S \in \mathcal{S}$,

$$C_\lambda(S) = \{x \in S : \text{ for all } y \in S,\ u(y) - u(x) \leq \theta(x)\}, \tag{3}$$

together with the regularity requirement:

$$u(x) \geq u(y) \implies u(x) + \theta(x) \geq u(y) + \theta(y). \tag{4}$$

Condition (4) is the coherence restriction that ensures the interval endpoints $u(x) + \theta(x)$ preserve the underlying ranking induced by u. It is satisfied by $\theta(\cdot)$ constructed from any monotone relabeling f as above.[2]

We now extend this idea into the family of correspondences $\{C_\lambda : \lambda \in \Lambda\}$. The intent is to keep a single benchmark utility u—interpretable as a stable preference ranking—and let cognition affect only the size of the tolerance bands around each alternative.

Definition 1 (Cognition-dependent representation (CDR))**.** The family $\mathcal{C} \coloneqq \{C_\lambda : \lambda \in \Lambda\}$ admits a *cognition-dependent representation* if and only if there exist functions $u : X \to \mathbb{R}$ and $\varepsilon : \Lambda \times X \to \mathbb{R}_+$ and a weak order $\geq$ on Λ with a maximal element $\overline{\lambda}$ such that, for all $\lambda \in \Lambda$ and all $S \in \mathcal{S}$,

$$C_\lambda(S) = \{x \in S : \forall y \in S, u(y) - u(x) \leq \varepsilon(\lambda, x)\},$$

with

(i) for all $x \in X$, $\varepsilon(\overline{\lambda}, x) = 0$,

[2]Technically we have shown only necessity of (3) and (4) for a semiordered representation. Sufficiency also holds, see Proposition 3 in the online appendix.

(ii) for all $x \in X$ and all $\lambda', \lambda \in \Lambda$,

$$\lambda' \geq \lambda \implies \varepsilon(\lambda', x) \leq \varepsilon(\lambda, x),$$

and

(iii) for all $x, y \in X$ and all $\lambda \in \Lambda$,

$$u(y) \geq u(x) \implies u(y) + \varepsilon(\lambda, y) \geq u(x) + \varepsilon(\lambda, x).$$

◊

Under Definition 1, x is admissible from S whenever it is close enough to the menu optimum in terms of the benchmark utility:

$$x \in C_\lambda(S) \iff \max_{y \in S} u(y) - u(x) \leq \varepsilon(\lambda, x).$$

Condition (ii) formalizes cognition as increasing discrimination: as cognition rises, the tolerated utility gap weakly shrinks for each x, so admissible sets weakly contract. Condition (i) states that at maximal cognition, tolerances vanish and the model reduces to exact maximization of u. This allows for the interpretation of u as a benchmark utility, as it represents the agents' "true" preferences, which would be apparent and actionable at maximal cognition.

Condition (iii) is the same regularity condition as in (4). It establishes that, if one fixes λ, each C_λ is semiorder-rationalizable: after a suitable increasing, λ-dependent rescaling of the utility axis, it would be possible to express admissibility using a constant gap. Doing so would however cause the utilities to become cognition dependent; as such, they could no longer be interpreted as a representation of benchmark preferences.

3 Choice axioms

We now ask: when does an observed family $\mathcal{C}$ admit a CDR?

3.1 WOCI

Definition 1 uses that cognition is weakly ordered: higher cognition corresponds to greater selectivity. In our axiomatic development, we impose a structure on the observable family $\{C_\lambda : \lambda \in \Lambda\}$ and a weak order can de derived from this structure.

Axiom WOCI (Weak Ordering of Cognition Indices)**.** For all $\lambda, \lambda' \in \Lambda$ and all $S, T \in \mathcal{S}$:

(i) **Nestedness.** Either $C_\lambda(S) \subseteq C_{\lambda'}(S)$ or $C_{\lambda'}(S) \subseteq C_\lambda(S)$.

(ii) **Directional consistency.** If there exists $S \in \mathcal{S}$ with $C_\lambda(S) \subsetneq C_{\lambda'}(S)$, then for every $T \in \mathcal{S}$, $C_\lambda(T) \subseteq C_{\lambda'}(T)$.

WOCI allows us to define an order of cognition levels purely from choice-set inclusion. Indeed, define a binary relation $\geq$ on Λ by

$$\lambda' \geq \lambda \iff C_{\lambda'}(S) \subseteq C_\lambda(S) \text{ for all } S \in \mathcal{S}.$$

Thus $\lambda' \geq \lambda$ means that λ' is (weakly) more selective than λ at every menu. Write $\lambda' > \lambda$ for the strict part: $\lambda' > \lambda$ if $\lambda' \geq \lambda$ and not $\lambda \geq \lambda'$. Under WOCI, the relation $\geq$ is a weak order on Λ (complete and transitive; see section A.1 in the appendix).

As an example, consider a person shopping for spaghetti. Under a low cognition level λ—perhaps because they are "in a hurry on a Monday morning"—the admissible set $C_\lambda(S)$ may be large: they eliminate only obvious non-starters and treat many options as good enough. Under a higher cognition index λ'—when they are "well-rested with no other commitments"—they apply finer distinctions and end up with a smaller admissible set $C_{\lambda'}(S)$. If for some menu S we observe $C_{\lambda'}(S) \subsetneq C_\lambda(S)$, then by directional consistency the same inclusion holds for every menu, so $\lambda' > \lambda$.

Because X is finite, $\mathcal{S}$ is finite, so there are only finitely many distinct correspondences mapping $\mathcal{S} \to \mathcal{S}$. Hence Λ can be partitioned into a finite number of $\geq$-equivalence classes, and the induced strict order on these classes has a minimum $\underline{\lambda}$ and maximum $\overline{\lambda}$. We henceforth let $\overline{\lambda}$ denote a member of the maximal class.

Definition 1 includes an (a priori unobserved) weak order over cognition labels, while Axiom WOCI constructs a weak order from observed choice. In the main characterization theorem (section 3.3), these coincide: a family $\{C_\lambda : \lambda \in \Lambda\}$ satisfies WOCI and WARPD if and only if it admits a cognition-dependent representation whose cognition order is exactly the weak order induced by WOCI.

3.2 WARP and WARPD

At maximal cognition, the cognition-dependent representation in Definition 1 imposes $\varepsilon(\overline{\lambda}, x) = 0$ for all x, so $C_{\overline{\lambda}}$ behaves like a standard utility-maximizing choice correspondence. Whenever this is the case, choices are said to be rational, and the canonical consistency condition for rationality is the Weak Axiom of Revealed Preference (WARP). We state it here first for a cognition-independent correspondence $C : \mathcal{S} \to \mathcal{S}$.

Axiom WARP (Weak Axiom of Revealed Preference)**.** For all $x, y \in X$ and all $S, T \in \mathcal{S}$ with $\{x, y\} \subseteq S \cap T$,

$$\text{If } x \in C(S) \text{ and } y \in C(T), \text{ then } x \in C(T) \text{ and } y \in C(S). \tag{5}$$

WARP captures the idea that choices reveal stable pairwise comparisons.In our setting, cognition matters and the admissible set may vary with λ, so imposing WARP separately at each cognition is too strong a condition. We instead impose a weaker requirement that anchors the revealed-preference comparison at an alternative that is cognitively robust.

Definition 2 (Robust admissibility)**.** For a menu S, say an alternative x is *robustly admissible in* S if it is admissible under every cognition index:

$$x \text{ is robust in } S \iff x \in \bigcap_{\lambda' \in \Lambda} C_{\lambda'}(S).$$

◊

Axiom WARPD (Weak Axiom of Revealed Preference with Difficulty)**.** For all $\lambda \in \Lambda$, all $x, y \in X$, and all $S, T \in \mathcal{S}$ with $\{x, y\} \subseteq S \cap T$,

$$x \in \bigcap_{\lambda' \in \Lambda} C_{\lambda'}(S) \text{ and } y \in C_\lambda(T) \implies x \in C_\lambda(T) \text{ and } y \in C_\lambda(S). \tag{6}$$

Axiom WARPD says: if x is robustly admissible in S while y is admissible from T at cognition λ, then at that same cognition level both x and y must be admissible in both menus. Intuitively, if an alternative survives in S no matter how selective the decision-maker is, it can be used as a point of comparison at lower cognition levels.

Under WOCI, robust admissibility is equivalent to admissibility at maximal cognition: since $C_{\overline{\lambda}}(S) \subseteq C_\lambda(S)$ for all λ, one has that

$$\bigcap_{\lambda' \in \Lambda} C_{\lambda'}(S) = C_{\overline{\lambda}}(S).$$

Using the WOCI reduction of robust admissibility, WARPD admits the following equivalent formulation.

Proposition 1. *Suppose Axiom WOCI holds. Then Axiom WARPD is equivalent to: for all $\lambda \in \Lambda$, all $x, y \in X$, and all $S, T \in \mathcal{S}$ with $\{x, y\} \subseteq S \cap T$,*

$$x \in C_{\overline{\lambda}}(S) \text{ and } y \in C_\lambda(T) \implies x \in C_\lambda(T) \text{ and } y \in C_\lambda(S). \tag{7}$$

In particular, setting $\lambda = \overline{\lambda}$ in (7) yields exactly WARP for the maximal-cognition correspondence $C_{\overline{\lambda}}$, while for $\lambda < \overline{\lambda}$ the WARP conclusion only holds when one of the two alternatives is selected at maximal cognition. This captures the idea of partially resolved preferences: while at $\overline{\lambda}$, the decision-maker acts according to their stable, "true" benchmark preferences, at any lower level $\lambda < \overline{\lambda}$ choices only approximate the benchmark. The intuition is as follows. Suppose x is the benchmark-optimal choice from S (so $x \in C_{\overline{\lambda}}(S)$), and, at a lower cognition level λ, an agent finds y admissibile from a menu T (where both menus contain x and y). For y to be admissible in T in the presence of the (benchmark) superior option x, it must be a "good enough" alternative at that level of cognition. WARPD requires this "good enough" status to be consistent in two ways. First, if y is good enough to be admissible in T alongside x, it must also be good enough to be admissible in S at that same λ. Second, since x is superior to y in the benchmark, if y is good enough for admissibility in T, then x has to also be admissible.

3.3 Equivalence

The main theorem relating cognition-dependent choice and utility can now be stated.

Theorem 1. *For a family of choice correspondences $\mathcal{C} = \{C_\lambda : \lambda \in \Lambda\}$, the following are equivalent:*

(i) $\mathcal{C}$ admits a cognition-dependent representation (Definition 1).

(ii) $\mathcal{C}$ satisfies WOCI and WARPD.

The proof that (ii) $\Longrightarrow$ (i) is constructive, and shows that the CDR is almost unique up to monotone transformations of u: once any represntation of the benchmark is selected, the thresholds ε are almost pinned down by the choice of u (subject to some wiggle room due to the finiteness of the problem). We also note that, under Theorem 1, the cognition monotonicity of Definition 1 is understood with respect to the weak order $\geq$ induced by WOCI.

4 Preference axioms

So far we have characterized cognition-indexed admissibility behavior $\mathcal{C} = \{C_\lambda : \lambda \in \Lambda\}$ in terms of WOCI and WARPD and showed their equivalence to the cognition-dependent representation (CDR). This section gives a preference-side formulation in terms of binary relations.

For each cognition index λ, we interpret an asymmetric relation P_λ as a *resolved strict comparison*: $xP_\lambda y$ means that, under cognition λ the agent can decisively act on the benchmark advantage of x over y. Its symmetric complement,

$$xI_\lambda y \iff x \not{P}_\lambda y \text{ and } y \not{P}_\lambda x,$$

is read as *unresolved* at that cognition level: neither strict statement is behaviorally actionable, so both alternatives remain co-admissible.

The structure we impose on $\mathcal{P} := \{P_\lambda : \lambda \in \Lambda\}$ is closely related to Roberts (1971). Although Roberts studies probabilistic consistency, his axioms on families of (deterministic) relations deliver a useful characterization of when certain indexed strict comparisons as *homogeneous nested family of semiorders* around a single, benchmark weak order. We show that this is exactly the preference-side content of the CDR.

4.1 Definitions

Given any binary relation P on X, define the induced (maximal-element) correspondence $\max_P : \mathcal{S} \to \mathcal{S}$ by

$$\max_P(S) := \{x \in S : \forall y \in S, y \not P x\}.$$

If P is acyclic, then $\max_P(S) \neq \varnothing$ for every nonempty menu S.

Definition 3 (Rationalizability by a relation)**.** A correspondence $C : \mathcal{S} \to \mathcal{S}$ is *rationalizable by P* if

$$C(S) = \max_P(S) \qquad \text{for all } S \in \mathcal{S}.$$

◊

Definition 4 ($\mathcal{P}$–rationalizability)**.** Let $\mathcal{C} = \{C_\lambda : \lambda \in \Lambda\}$ be a family of nonempty–valued choice correspondences. We say that $\mathcal{C}$ is *$\mathcal{P}$–rationalizable* by a family of relations $\mathcal{P} = \{P_\lambda : \lambda \in \Lambda\}$ if, for every $\lambda \in \Lambda$ and every menu $S \in \mathcal{S}$,

$$C_\lambda(S) = \max_{P_\lambda}(S).$$

◊

A semiorder is the classic strict-comparison structure behind just-noticeable differences: local unresolvedness (I) need not be transitive, so chains of local indistinguishabilities do not imply global indistinguishability.

Definition 5 (Semiorder)**.** An asymmetric binary relation P on X is a *semiorder* if it satisfies:

(i) **Interval-order condition.** For all $x, y, z, w \in X$, xPy and zPw imply xPw or zPy.

(ii) **Semi-transitivity.** For all $x, y, z, w \in X$, xPy and yPz imply xPw or wPz.

◊

As recalled elsewhere in the paper (and classical in the measurement literature), a semiorder admits a constant-threshold representation (Scott and Suppes (1958); see also Luce, 1956, Fishburn, 1985): there exist $v : X \to \mathbb{R}$ and $\theta \geq 0$ such that $xPy \iff v(x) > v(y) + \theta$.

The next definition is central in Roberts (1971). It links a semiorder (resolved strict comparisons) to a weak order (benchmark ranking) by requiring that unresolvedness be "convex" along the benchmark.

Definition 6 (Compatibility Roberts, 1971)**.** Let P be an asymmetric relation on X and let $\succeq$ be a weak order on X. Let I be the symmetric complement of P. We say that P is *compatible with* $\succeq$ if, for all $x, y, z \in X$,

(i) $xPy \Rightarrow x \succeq y$.

(ii) If $x \succsim y \succsim z$ and xIz, then xIy and yIz.

◊

Condition (ii) says: if x and z are unresolved, then every benchmark-intermediate alternative y must also be unresolved with each endpoint. In words, the I-band is an interval in the benchmark order.

Definition 7 (Homogeneous family of semiorders)**.** A family $\mathcal{P} = \{P_\lambda : \lambda \in \Lambda\}$ of asymmetric relations on X is a *homogeneous family of semiorders* if there exists a weak order $\succsim$ on X such that every P_λ is compatible with $\succsim$. ◊

By Roberts (1971, Theorem 1), an asymmetric relation is a semiorder if and only if it is compatible with some weak order. Hence Definition 7 indeed describes a family of semiorders sharing a common benchmark weak order.

4.2 Roberts' axioms

Fix a family $\mathcal{P} := \{P_\lambda : \lambda \in \Lambda\}$ of asymmetric relations on X. The following are Axioms 1–3 of Roberts (1971), rewritten in our notation.

Axiom GA (General asymmetry)**.** For all $x, y \in X$ and all $\lambda, \lambda' \in \Lambda$,

$$xP_\lambda y \implies y \not{P}_{\lambda'} x.$$

GA rules out reversals across cognition regimes: once a strict comparison is ever behaviorally resolved in one direction, it is never resolved in the opposite direction under any regime.

Axiom N (Nestedness)**.** For all $\lambda, \lambda' \in \Lambda$,

$$P_\lambda \subseteq P'_\lambda \quad \text{or} \quad P'_\lambda \subseteq P_\lambda.$$

N is the preference-side counterpart of WOCI: higher cognition resolves (weakly) more strict comparisons, hence yields a (weakly) larger strict relation.

Axiom UW (Uniform witnessing)**.** For all $x, y \in X$, if $xI_\lambda y$ for all $\lambda \in \Lambda$, then for all $z \in X$ and all $\lambda \in \Lambda$,

$$(xP_\lambda z \iff yP_\lambda z) \text{ and } (zP_\lambda x \iff zP_\lambda y).$$

UW states that if two alternatives are never behaviorally distinguished at any cognition level, then they are indistinguishable from the standpoint of all other resolved comparisons.

Define a not-worse-than relation $\succsim$ by

$$x \succsim y \iff y \not{P}_\lambda x \text{ for all } \lambda \in \Lambda. \tag{8}$$

The remaining axioms below are alternative ways to enforce that $\succeq$ is transitive and that the family is compatible with it. Roberts shows they are equivalent given GA, N and UW.

Axiom OD (Ordered discovery)**.** For all $x, y, z \in X$ with $x \succeq y$ and $y \succeq z$:

(i) $x \succeq z$.

(ii) For every $\lambda \in \Lambda$, $xP_\lambda y$ or $yP_\lambda z$ implies $xP_\lambda z$.

OD is Roberts's deterministic analogue of "strong stochastic transitivity": if x is benchmark-above y and y benchmark-above z, then any resolved strict comparison involving an adjacent pair forces the outer comparison to be resolved as well.

Axiom SSC (Second semiorder condition)**.** For all $x, y, z, w \in X$ and all $\lambda, \lambda' \in \Lambda$,

$$xP_\lambda y \text{ and } yP_{\lambda'} z \implies xP_\lambda w \text{ or } wP_{\lambda'} z.$$

SSC is a generalization of semitransitivity. It rules out a having $xP_\lambda yP_{\lambda'} z$ together with a "shield" w that simultaneously escapes defeat by x at level λ and also escapes defeating z at level λ'. Thus, whenever x beats y at λ and y beats z at λ', every alternative w must lie on one side or the other of the "gap" between x and z: either w is sufficiently below x to be beaten by x at λ, or w is sufficiently above z to beat z at λ'.

Axiom WSSC (Weak second semiorder condition)**.** For all $x, y, z \in X$ and all $\lambda, \lambda' \in \Lambda$,

$$xP_\lambda y \text{ and } yP_{\lambda'} z \implies xP_\lambda z \text{ or } xP_{\lambda'} z.$$

Fix λ, λ'. If $xP_\lambda y$ and $yP_{\lambda'} z$, then WSSC requires that the composite comparison between x and z is resolved at one of the two levels: either λ already resolves x against z or λ' does. Any time a strict chain is witnessed using λ for the first link and λ' for the second, the endpoint comparison must be detectable without needing a third level. In particular, WSSC says that resolved comparisons do not require "accumulating" cognition levels along a path: one of the indices that resolves a link in the chain must already be strong enough to resolve the endpoints.

4.3 Second Equivalence

We first note two consequences of finiteness and GA, N. First, since X is finite, there are only finitely many distinct binary relations on X. By N, that means that we can weakly order cognition by inclusion: $\lambda \geq \lambda' \iff P_\lambda \supseteq P_{\lambda'}$. And, as we did with WOCI, we conclude that there are finitely many equivalence classes of this weak order, and we can take $\bar{\lambda}$ as a representative member of the maximal class.[3]

[3] See Lemma 2 in the Appendix for details.

Second, note that $\succeq$ defined in (8) is complete and reflexive. Moreover, if N holds, then[4]

$$x \succeq y \iff y \not{P}_{\lambda} x.$$

The next statement is a direct specialization of Roberts (1971, Theorem 5) to our setting.

Proposition 2 (Roberts, 1971)**.** *Let X be finite and let $\mathcal{P} = \{P_\lambda : \lambda \in \Lambda\}$ satisfy GA, N, and UW. Then the following are equivalent:*

(i) OD holds.

(ii) SSC holds.

(iii) WSSC holds.

(iv) $\mathcal{P}$ is a homogeneous family of semiorders with benchmark $\succeq$.

Proof. See the proof of Roberts (1971, Theorem 5). ■

As it turns out, Axioms GA, N and UW, together with either one of conditions explicited in the Proposition, is exactly the necessary and sufficient conditions for a $\mathcal{P}$-rationalizable $\mathcal{C}$ to admit a CDR.

Theorem 2. *For a family of choice correspondences $\mathcal{C} = \{C_\lambda : \lambda \in \Lambda\}$, the following are equivalent:*

(i) $\mathcal{C}$ admits a cognition-dependent representation in the sense of Definition 1.

(ii) $\mathcal{C}$ is $\mathcal{P}$–rationalizable (Definition 4) by a family $\mathcal{P} = \{P_\lambda : \lambda \in \Lambda\}$ satisfying GA, N, UW, and one (hence all) of OD, SSC, WSSC.

(iii) $\mathcal{C}$ is $\mathcal{P}$–rationalizable (Definition 4) by a homogeneous family of semiorders (Definition 7) $\mathcal{P} = \{P_\lambda : \lambda \in \Lambda\}$ satisfying GA, N, UW.

Since choices from all menus are observed, whenever $\mathcal{C}$ is $\mathcal{P}$-rationalizable as in (ii) and (iii), we have that this rationalizing $\mathcal{P}$ can be identified from choices from binary menus.

Combining Theorem 2 with Theorem 1 also yields the bridge WOCI and WARPD and rationalizability by a homogeneous family of semiorders (or any one of the other equivalent properties in Proposition 2). In this representation, the benchmark weak order is stable and represent rational preferences in the classical sense, while cognition governs which strict comparisons (hence which defeats) are effectively acted upon.

[4]See Lemma 3 in the Appendix for details.

5 Fuzzy axioms

In Sections 2–4 we analyzed cognition-indexed choice behavior through a family of admissibility correspondences $\mathcal{C} = \{C_\lambda : \lambda \in \Lambda\}$ and provided three equivalent characterizations: (i) the cognition-dependent threshold representation (Definition 1), (ii) the choice axioms WOCI and WARPD, and (iii) the preference-side axioms GA, N, UW for the homogeneous family of semiorders $\{P_\lambda : \lambda \in \Lambda\}$.

This section provides a further, conceptually distinct representation: we show that the entire cognition-indexed system can be encoded by a *single fuzzy relation* that assigns to each ordered pair (x, y) an ordinal degree of *comparison difficulty* for recognizing $y \succ x$. The observed cognition levels then correspond to cuts of this fuzzy relation under an appropriate calibration. Besides offering intuition, the fuzzy formulation clarifies a link between two classic responses to the Sorites paradox: the just-noticeable-difference (JND) account formalized by semiorders and the graded-truth account formalized by fuzzy sets (See section 5.5).

5.1 Fuzzy sets, cuts, and ordinal meaning

Unlike classical set theory, where an element either belongs to a set or it does not, fuzzy set theory introduces the notion of partial degrees of membership. This approach allows for the representation of concepts that are inherently vague or imprecise in a structured, quantitative manner. The degree of membership for each element is typically represented by a value from a totally ordered set V, known as a truth set. This set is bounded by a smallest element, $\underline{v}$, representing complete non-membership, and a greatest element, $\bar{v}$, representing full membership. Intermediate values capture varying degrees of partial membership.

As an illustration, let X be a population of individuals and consider the predicate "tall." In natural language, "tall" is vague: people broadly agree on its meaning but not on a precise cutoff. A fuzzy set captures this by assigning to each height h a membership value $\mu_{\text{Tall}}(h) \in V$. For example, $\mu_{\text{Tall}}(170\text{cm}) \in (\underline{v}, \bar{v})$ reflects partial agreement that 170 cm (about $5'7''$) counts as tall.

Membership degrees admit well-known interpretations beyond graded inclusion. One view treats $\mu_A(x)$ as a *similarity* score between x and a prototype (or ideal) of A. Another interprets $\mu_A(x)$ as an ordinal assessment of *plausibility*, *belief*, or graded *truth* in the sense of fuzzy logic (Zadeh, 1971, Liu and Kerre, 1998, Shimoda, 2002). These readings are conceptually distinct from probabilities: typically only the *order* of the values matters, not their cardinal scale.

Formally, fix a totally ordered set V (a *truth set*) with bottom element $\underline{v}$ and top element $\bar{v}$. A *fuzzy set* on X is a map $\mu : X \to V$; the value $\mu(x)$ is the degree to which the statement "x belongs to the set" is deemed true or valid in a given context.

A standard way to connect fuzziness to crisp sets is by *cuts*. Given μ and a threshold

$v \in V$, the v-cut is the crisp set $\{x : \mu(x) \geq v\}$. As v rises, the cut becomes smaller: higher thresholds correspond to stricter admission criteria. For instance, if X is a population and $\mu_{\text{Tall}}(x)$ assigns each person x a degree of tallness, then choosing a threshold $v \in V$ produces the crisp set of individuals whose tallness is at least v: $\{x \in X : \mu_{\text{Tall}}(x) \geq v\}$. If we raise the threshold from v to $v' > v$, fewer people qualify as tall—someone with $\mu_{\text{Tall}}(x) \in [v.v')$ is included at level v but excluded at level v'.

5.1.1 Fuzzy binary relations

Fuzzy sets extend naturally to binary relations. A (binary) *fuzzy relation* on X is a map

$$R : X \times X \to V.$$

The value $R(x, y)$ is the membership degree of the ordered pair (x, y) in the relation's fuzzy graph— that is, the extent to which the statement "x is R-related to y" is deemed true. Note that classical (crisp) relations are the special case where V has only two elements, $\{\underline{v}, \bar{v}\}$, or $\{\text{True}, \text{False}\}$, and R takes only these extreme values.

As with fuzzy sets, intermediate truth values in a fuzzy relation admit multiple, well-known interpretations. For a pair (x, y), the value $R(x, y) \in (\underline{v}, \bar{v})$ may be read as the *intensity* with which the relation holds, as a *similarity* score between x and y,or as an epistemic grade of belief or confidence that the statement "x is R-related to y" is true.[5]

In this paper, we interpret $R(x, y)$ as an ordinal measure of how difficult it is to recognize a rational-benchmark preference "$y \succ x$." That is, when $R(x, y)$ is small, we take it to mean that it is easy to tell that y beats x. On the other hand, if $R(x, y)$ is large, it takes a higher degree of discrimination to conclude that y beats x. Finally, if $R(x, y) = \bar{v}$, recognizing $y \succ x$ is impossible (i.e. $y \succ x$ is false in the rational-benchmark order, implying $x \succsim y$).

This means that, in this paper, mistakes are omissions, not commissions. The agent may fail to perceive a true strict preference if they are constrained in their ability to discriminate, but they never "discover" a strict preference that contradicts their own benchmark ranking.

These observations are captured by the fuzzy analogue of the completeness property for crisp binary relations:

Definition 8 (Fuzzy completeness)**.** A fuzzy relation R is *fuzzy complete* if for all $x, y \in X$,

$$\max\{R(x, y), R(y, x)\} = \bar{v}.$$

◊

Fuzzy completeness says that, for every pair (x, y), at least one of the two strict statements $y \succ x$ and $x \succ y$ is impossible to recognize, because at least one of them is false

[5] These are almost always ordinal concepts: only the ranking of values in V is meaningful, not their cardinal magnitudes.

in the benchmark weak order.[6]

We now impose the same intuition presented before for ordered discovery: if y is benchmark-preferred to x, then recognizing $y \succ x$ should not be *harder* than recognizing a smaller benchmark-ranking separation. Equivalently, comparisons between benchmark-far-apart alternatives should be discoverable (weakly) earlier.

A technically convenient way to impose exactly this monotonicity, while also ensuring that the benchmark weak order is transitive, is the following fuzzy transitivity condition, which coincides with the weakest t-norm-based transitivity used in the fuzzy-relations literature (the "drastic" t-norm). See Ovchinnikov (1981) for background.

Definition 9 (Fuzzy transitivity)**.** A fuzzy relation R is *fuzzy transitive* if for all $x, y, z \in X$,

$$\max\{R(x,y), R(y,z)\} = \bar{v} \implies R(x,z) \geq \min\{R(x,y), R(y,z)\}.$$

◊

Definition 10 (Fuzzy rationality)**.** A fuzzy relation R is *fuzzy rational* if it satisfies fuzzy completeness (Definition 8), and fuzzy transitivity (Definition 9). ◊

Under fuzzy rationality, the benchmark weak order can be read off from the $\bar{v}$-entries:

$$x \succeq y \quad \Longleftrightarrow \quad R(y,x) = \bar{v}.$$

This is exactly the benchmark order that appears at maximal cognition in our model. Note that fuzzy completeness and fuzzy transitivity imply $\succsim$ is complete and transitive, respectively, and thus can be properly interpreted as a rational benchmark.

5.2 From a fuzzy relation to crisp choice

A fuzzy relation produces a *nested* family of crisp strict comparisons and thus a nested family of crisp choice correspondences, by cutting at different resolution thresholds.

To see how, fix $v \in V$. Define a crisp strict relation P^v by

$$yP^v x \quad \Longleftrightarrow \quad R(x,y) < v. \tag{9}$$

Thus $yP^v x$ means: at resolution level v, the decision maker can recognize that y beats x. If $v' \geq v$, then $P^v \subseteq P^{v'}$: higher resolution reveals more strict comparisons.

Given R, define the fuzzy "choice score" of x in menu S by

$$C^R(S,x) := \min_{y \in S} R(x,y) \qquad (x \in S). \tag{10}$$

This is the difficulty of the easiest recognized defeat of x within S: if some y beats x very easily, then $C^R(S,x)$ is small.

[6] If x and y are benchmark-indifferent, then both directions are impossible, so both values equal $\bar{v}$.

Given $v \in V$, the corresponding crisp choice correspondence is the v-cut

$$C^{R;v}(S) := \{x \in S : C^R(S,x) \geq v\}. \tag{11}$$

Equivalently, $x \in C^{R;v}(S)$ iff there is no $y \in S$ such that $R(x,y) < v$, that is, if and only if x is undominated by P^v. Hence

$$C^{R;v}(S) = \max_{P^v}(S). \tag{12}$$

Note that, as v rises, $C^{R;v}(S)$ shrinks.

5.3 Fuzzy rationalizability

The preceding construction produces a family of correspondences indexed by $v \in V$. Our data produce a family indexed by cognition $\lambda \in \Lambda$. To connect them, we introduce a calibration ν, mapping cognition into resolution thresholds.

Definition 11 (Calibration)**.** Let Λ be a weakly ordered set of cognition indices with a maximal element $\overline{\lambda}$ and let V be a totally ordered truth set with top element $\bar{v}$. A map $\nu : \Lambda \to V$ is a *calibration* if it is order-preserving and normalized at the top:

$$\lambda' \geq_\Lambda \lambda \implies \nu(\lambda') \geq_V \nu(\lambda), \qquad \nu(\overline{\lambda}) = \bar{v}.$$

◊

Definition 12 (Fuzzy regular-rationalizability)**.** A family of correspondences $\mathcal{C} = \{C_\lambda : \lambda \in \Lambda\}$ is *fuzzy regular-rationalizable* if there exist a truth set V, a fuzzy-rational relation $R : X \times X \to V$, a weak order $\geq$ on Λ with maximal element $\overline{\lambda}$, and a calibration $\nu : \Lambda \to V$ such that for all $\lambda \in \Lambda$ and $S \in \mathcal{S}$,

$$C_\lambda(S) = C^{R;\nu(\lambda)}(S).$$

◊

The interpretation is direct: $\nu(\lambda)$ is the resolution deployed by the decision maker in cognition level λ. Higher cognition means a higher threshold, which reveals more strict comparisons and therefore yields a smaller set of admissible options.

If $\mathcal{C}$ admits a cognition-dependent representation (Definition 1), then fuzzy regular-rationalizability can be proven. The next proposition makes this explicit and is the key step behind the equivalence theorem below.

Proposition 3. *Suppose $\mathcal{C}$ admits a cognition-dependent representation with benchmark utility u and threshold function $\varepsilon(\lambda, x)$ (Definition 1). Let V be any totally ordered set and let $\nu : \Lambda \to V$ be any calibration with $\nu(\overline{\lambda}) = \bar{v}$.*

Define $R : X \times X \to V$ *by*[7]

$$R(x,y) := \sup\{\nu(\lambda) : u(y) - u(x) \leq \varepsilon(\lambda, x)\}. \tag{13}$$

and, if $u(y) - u(x) > \varepsilon(\lambda, x)$ *for all* λ*, let* $R(x,y) = \underline{v}$*.*

Then for every $\lambda \in \Lambda$ *and every menu* $S \in \mathcal{S}$*,*

$$C_\lambda(S) = C^{R;\nu(\lambda)}(S).$$

The meaning of (13) is simple: $R(x,y)$ is the largest calibrated resolution level at which x still counts as "good enough" relative to y. Once cognition exceeds that threshold, y becomes a recognized strict improvement over x and can eliminate it.

5.4 Third Equivalence

We can now state the final equivalence.

Theorem 3. *For a family of choice correspondences* $\mathcal{C} = \{C_\lambda : \lambda \in \Lambda\}$*, the following are equivalent:*

(i) $\mathcal{C}$ *admits a cognition-dependent representation (Definition 1).*

(ii) $\mathcal{C}$ *is fuzzy regular-rationalizable. (Definition 12).*

Proposition 3 yields (i) $\Longrightarrow$ (ii). Conversely, (ii) $\Longrightarrow$ (i) follows because cut correspondences are automatically nested (hence satisfy WOCI once Λ is ordered by the calibration) and fuzzy rationality of R enforces the consistency encoded by WARPD. We then apply Theorem 1.

Under fuzzy rationalizability, the strict relations revealed from binary menus,

$$yP_\lambda x \iff x \notin C_\lambda(\{x,y\}),$$

coincide with the calibrated cut relations $P^{\nu(\lambda)}$ induced by R in (9). Hence the fuzzy object R can be read as a compact encoding of the nested family $\{P_\lambda : \lambda \in \Lambda\}$: it assigns to each true strict benchmark comparison the cognition threshold at which that comparison becomes detectable.

5.5 The Sorites paradox

The classic Sorites paradox arises from the tension between (i) the apparent *tolerance* of vague predicates to small changes and (ii) the existence of clear endpoint cases. The paradox begins with a vague concept such as "heap", "bald", or "tall" and a sequence of objects that

[7]Note that, because X is finite, we can avoid an extra completeness assumption. Indeed, calibrations are constant on each WOCI-equivalence class.

differ only by imperceptibly small increments. One then accepts a local principle: a single tiny change should not flip the truth of the predicate (removing one grain should not turn a heap into a non-heap, removing one hair should not turn a non-bald person into bald, adding one gram of sugar should not turn something from not-sweet to sweet). Iterating this step along the whole sequence yields an absurd conclusion: if no single grain matters, then even removing thousands of grains cannot matter, so a heap never ceases to be a heap; if no single hair matters, then someone with one hair is not bald; if no single gram matters, then even a large sugar difference should not change sweetness. The paradox is that each local step is plausible in isolation, yet their unrestricted transitive closure conflicts with the evident fact that sufficiently many small changes *do* alter classification.

In discrimination models, the analogous fallacy is to treat local indistinguishability as transitive. Semiorders provide a standard way of getting around the paradox: nearby alternatives can be mutually incomparable (within the tolerance band) even though far-apart alternatives are strictly ranked. Thus the relation

$$xI_\lambda y \iff \quad \{x,y\} \subseteq C_\lambda(\{x,y\})$$

need not be transitive. One can have $x_0 I_\lambda x_1, x_1 I_\lambda x_2, \ldots, x_{n-1} I_\lambda x_n$ while $x_n P_\lambda x_0$ at the same cognition level. In the heap example, this corresponds to a sequence of piles that are pairwise indistinguishable when they differ by one grain, yet sufficiently separated piles are distinguishable.

Fuzzy set theory resolves Sorites in a seemingly different way. It rejects the idea that vague predicates must be either fully true or fully false. Instead, truth can change gradually. For example, along a sequence of heights, the statement "tall" can move smoothly from low to high membership. The paradoxical step that treats each local change as preserving full truth is blocked.

What we show is that, under fuzzy rationalizability, these two resolutions are in fact the same structure expressed in two languages. Fix λ and write $v := \nu(\lambda)$. Indeed,

$$xI_\lambda y \iff R(x,y) \geq v \text{ and } R(y,x) \geq v.$$

The left-hand side mean the semiordered indistinguishability; the right-hand side means "$x \succeq y$" and "$y \succeq x$" are "sufficiently true" (considering a "truth" threshold of v).

In a Sorites chain of local indistinguishabilities, say, $xI_\lambda y$ and $yI_\lambda z$, it remains possible that the endpoints are strictly comparable; for instance $xP_\lambda z$, which in the cut language is

$$R(z,x) < v,$$

which fuzzy transitivity explicitly allows. If we read $R(\cdot,\cdot)$ as graded support for "z preferred to x", then the statement is not true enough" (again, considering a particular "truth" threshold of v).

6 Summary Theorem

For convenience we collect all of the equivalences in a single Theorem.

Theorem 4. *For a family of choice correspondences $\mathcal{C} = \{C_\lambda : \lambda \in \Lambda\}$, the following are equivalent:*

(i) $\mathcal{C}$ *admits a cognition-dependent representation (Definition 1).*

(ii) $\mathcal{C}$ *satisfies WOCI and WARPD.*

(iii) $\mathcal{C}$ *is* $\mathcal{P}$*–rationalizable (Definition 4) by some family* $\mathcal{P} = \{P_\lambda : \lambda \in \Lambda\}$ *satisfying GA, N, UW, and one (hence all) of OD, SSC, WSSC.*

(iv) $\mathcal{C}$ *is* $\mathcal{P}$*–rationalizable (Definition 4) by some homogeneous family of semiorders (Definition 7)* $\mathcal{P} = \{P_\lambda : \lambda \in \Lambda\}$ *satisfying GA, N, UW.*

(v) $\mathcal{C}$ *is fuzzy regular-rationalizable. (Definition 12).*

7 Extension: Incompleteness of observation

The baseline theory treats cognition-indexed behavior as observed over a set of cognition levels Λ that contains a maximal (benchmark) cognition level $\overline{\lambda}$. When $\overline{\lambda}$ is observed, the benchmark ranking can be read directly from $C_{\overline{\lambda}}$ (or from $P_{\overline{\lambda}}$), and all representation results apply after restricting attention to the observed family.

In applications, however, one may only observe choices under a *subset* $\tilde{\Lambda} \subseteq \Lambda$ of cognition regimes. This is innocuous when $\overline{\lambda} \in \tilde{\Lambda}$, but it becomes more subtle when the benchmark level is *missing*. In that case, the full benchmark strict order $P_{\overline{\lambda}}$ is no longer directly observable. Nevertheless, our comparative-statics structure still implies that some benchmark comparisons are *already forced* by the observed levels. This subsection provides an explicit construction of those forced comparisons.

Throughout, we shall work with (crisp) strict preference relations. For each observed $\lambda \in \tilde{\Lambda}$, let P_λ be the strict relation recovered from binary menus (Section 4.1):

$$xP_\lambda y \iff y \notin C_\lambda(\{x, y\}).$$

Equivalently, the analyst may treat $\{P_\lambda"\lambda \in \tilde{\Lambda}\}$ as the primitive observed object whenever binary comparisons are elicited directly.

We assume that the *full* (possibly unobserved) family $\mathcal{P} = \{P_\lambda : \lambda \in \Lambda\}$ satisfies the preference axioms of Section 4: nestedness of strict comparisons and general asymmetry (GA and N), uniform witness (UW and ordered discovery OD. In particular, we assume there exists a maximal cognition level $\overline{\lambda} \in \Lambda$ such that $P_\lambda \subseteq P_{\overline{\lambda}}$ for all λ, and $P_{\overline{\lambda}}$ is the strict part of a benchmark weak order $\succeq$.

7.1 A revealed lower bound on the unobserved benchmark

The construction below uses *only* the observed relations $\{P_\lambda : \lambda \in \tilde{\Lambda}\}$. Intuitively, $xT_{\tilde{\Lambda}}y$ means that, across all observed cognition levels, x weakly dominates y in the sense that x inherits y's victories and is not more vulnerable than y to defeats, and that this dominance is strict at least somewhere.

Proposition 4. *Let $\tilde{\Lambda} \subseteq \Lambda$ be nonempty. Define a binary relation $T_{\tilde{\Lambda}}$ on X by, for $x, y \in X$,*

$$\begin{aligned} xT_{\tilde{\Lambda}}y \iff & \big(\forall \lambda \in \tilde{\Lambda}, \forall z \in X : [yP_\lambda z \implies xP_\lambda z] \text{ and } [zP_\lambda x \implies zP_\lambda y]\big) \\ & \text{and} \\ & \big(\exists \lambda_0 \in \tilde{\Lambda}, \exists z \in X : [xP_{\lambda_0} z \text{ and } y\not P_{\lambda_0} z] \text{ or } [zP_{\lambda_0} y \text{ and } z\not P_{\lambda_0} x]\big). \end{aligned}$$

Then, for every $\lambda \in \tilde{\Lambda}$,

$$P_\lambda \subseteq T_{\tilde{\Lambda}} \subseteq P_{\overline{\lambda}}.$$

In particular, if $\overline{\lambda} \in \tilde{\Lambda}$ then $T_{\tilde{\Lambda}} = P_{\overline{\lambda}}$.

Proposition 4 can be read as a partial-identification result: even when the benchmark level $\overline{\lambda}$ is not observed, the relation $T_{\tilde{\Lambda}}$ is computable from the observed cognition levels alone and is guaranteed (under the model) to be a subset of the benchmark strict preference $P_{\overline{\lambda}}$. Moreover, $T_{\tilde{\Lambda}}$ contains every strict comparison that appears at any observed level.

7.2 Monotonicity in observations

The next corollary clarifies how $T_{\tilde{\Lambda}}$ varies with the amount of observed cognition information. Importantly, the monotonicity is with respect to *set inclusion* of observed levels (more observed regimes), not with respect to the cognition order itself.

Corollary 1. *Let $\Lambda' \subseteq \Lambda'' \subseteq \Lambda$ be nonempty sets of observed cognition indices. Then*

$$T_{\Lambda'} \subseteq T_{\Lambda''}.$$

In particular, enlarging the set of observed cognition indices can only expand the revealed relation $T_{\tilde{\Lambda}}$.

Corollary 1 does *not* imply that observing a single higher cognition level is always more informative than observing a single lower cognition level. The next example illustrates that $T_{\{\lambda\}}$ can change non-monotonically with λ.

Example 2. Let $X = \{x, y, z\}$ and consider two levels $\lambda' < \lambda'' < \overline{\lambda}$. Let

$$P_{\lambda'} = \{(z, y)\}, \qquad P_{\lambda''} = \{(z, y), (z, x)\},$$

so $P_{\lambda'} \subseteq P_{\lambda''}$.

At level λ', the pair (x, y) belongs to $T_{\{\lambda'\}}$: the universal dominance clause is vacuous (since y has no outgoing edges and x has no incoming edges), and the witness clause holds because $zP_{\lambda'}y$ while $z\not{P}_{\lambda'}x$.

At level λ'', the pair (x, y) fails to belong to $T_{\{\lambda''\}}$: the only candidate witness z satisfies $zP_{\lambda''}y$ *and* $zP_{\lambda''}x$, so the witness clause fails.

Thus $xT_{\{\lambda'\}}y$ but $x\not{T}_{\{\lambda''\}}y$, even though $\lambda'' > \lambda'$. $\Diamond$

The point of Example 2 is that $T_{\tilde{\Lambda}}$ is monotone in *how many* cognition levels we observe (Corollary 1), but $T_{\{\lambda\}}$ need not be monotone in the cognition order when we replace one observed regime by another.

7.3 The $T_{\tilde{\Lambda}}$ relation and transitive cores

The definition of $T_{\tilde{\Lambda}}$ has a natural graph-theoretic interpretation: the first clause says that at each observed level λ, the set of alternatives beaten by y is contained in the set beaten by x, and the set of alternatives that beat x is contained in the set that beat y. This is exactly the contour-inclusion notion that appears in Nishimura (2018)'s *transitive core*.

To make this connection precise, associate to each λ the natural weak companion of P_λ,

$$\succeq_\lambda^{\text{wk}} := \{(a, b) \in X \times X : b\not{P}_\lambda a\}.$$

Equivalently, $\succeq_\lambda^{\text{wk}} = P_\lambda \cup I_\lambda$, where I_λ is incomparability induced by P_λ. The relation $\succeq_\lambda^{\text{wk}}$ is reflexive and complete, but need not be transitive.

Given a binary relation R on X, let $\uparrow_R(a) := \{b \in X : aRb\}$ and $\downarrow_R(a) := \{b \in X : bRa\}$ be its upper and lower contour sets. Following Nishimura (2018), define the *transitive core* of R by

$$a\,\mathrm{TC}(R)\,b \iff \uparrow_R(b) \subseteq \uparrow_R(a) \text{ and } \downarrow_R(a) \subseteq \downarrow_R(b).$$

The relation $\mathrm{TC}(R)$ is a preorder, and it is the largest subrelation $S \subseteq R$ such that R is S-transitive (see Nishimura, 2018).

Lemma 1. *For each $\lambda \in \Lambda$ and all $x, y \in X$,*

$$x\,\mathrm{TC}\big(\succeq_\lambda^{\text{wk}}\big)\,y \iff \big(\forall z \in X : yP_\lambda z \implies xP_\lambda z\big) \text{ and } \big(\forall z \in X : zP_\lambda x \implies zP_\lambda y\big).$$

Given $\tilde{\Lambda} \subseteq \Lambda$, define the preorder

$$\succsim_{\tilde{\Lambda}} := \bigcap_{\lambda \in \tilde{\Lambda}} \mathrm{TC}\big(\succeq_\lambda^{\text{wk}}\big),$$

and write $\succ_{\tilde{\Lambda}}$ for its strict (asymmetric) part.

Proposition 5. *For every nonempty $\tilde{\Lambda} \subseteq \Lambda$,*

$$T_{\tilde{\Lambda}} = \succ_{\tilde{\Lambda}} = \{(x, y) : x \succsim_{\tilde{\Lambda}} y \text{ and } y \not\succsim_{\tilde{\Lambda}} x\}.$$

In particular, $T_{\tilde{\Lambda}}$ is transitive and irreflexive (a strict partial order).

Proposition 5 shows that $T_{\tilde{\Lambda}}$ extracts the *common transitive content* of the observed weak relations $\{\succeq_{\lambda}^{\mathrm{wk}}\}_{\lambda\in\tilde{\Lambda}}$. Combined with Proposition 4, it follows that when the data are generated by our cognition model, this common transitive content is guaranteed to be a subset of the latent benchmark strict order $P_{\bar{\lambda}}$.

8 Related literature

This paper contributes to several strands of literature on (i) semiorders and threshold models of imperfect discrimination, (ii) deterministic set-valued choice under attention, procedural and cognitive constraints, (iii) decisions under information-processing limiations, and finally (iv) fuzzy set theory as a formal account of vagueness and graded comparison.

8.1 Semiorders and just-noticeable differences

Semiorders grew out of attempts to formalize the psychophysical idea that very small differences are not reliably perceived. In the Weber–Fechner tradition, discrimination is limited by a just-noticeable difference (JND), so local indistinguishability should not be treated as globally transitive (Weber, 1834, Fechner, 1860). Luce's motivating examples (e.g. a sequence of coffee cups with incrementally more sugar) make the point vivid: one may fail to discriminate adjacent stimuli while still sharply preferring sufficiently separated endpoints, so the usual move of modeling indifference as a transitive equivalence relation becomes untenable (Luce, 1956). Semiorders formalize this by allowing a thick indifference band, typically implemented by a constant threshold in a numerical representation. The foundational representation result is the Scott–Suppes theorem: on finite domains, a semiorder admits a utility representation with a fixed threshold (often normalized to 1), which in turn implies a strong restriction on admissible rescalings of the representing index relative to weak orders (uniqueness essentially up to positive affine transformations) (Scott and Suppes, 1958). Fishburn is credited with many of the subsequent developments in the theory of semiorders and interval orders (Fishburn, 1985, 1970a). A related theme concerns measurement: because the semiorder threshold is invariant only under affine changes of scale, semiorders sit between purely ordinal rankings and genuinely cardinal measurement, and this intermediate status has motivated careful analyses of representation and "measurement paradox"-type issues for imperfect or partial preference structures (Beja and Gilboa, 1992).

A second strand asks when semiorder behavior can be recovered from observable choice. Jamison and Lau give a revealed-preference axiomatization for semiorder-rationalizable choice, and later work refines and streamlines such conditions on standard finite-menu domains (Jamison and Lau, 1973, Fishburn, 1970b). More recent contributions continue to deploy semiorders as primitives for modeling coarse comparisons, sequential evaluation, or approximate rationality. Manzini and Mariotti's lexicographic semiorders combine a semiorder thresholds with lexicographic priority. Dziewulski (2020) connects JND-style discrimination directly to revealed-preference measures of approximate utility maximization. Balart (2021) uses semiorder preferences in a Hotelling setting. Against this background, our paper's contribution is to move from the classic "single correspondence" problem to a *linked family* C_λ in which discrimination varies with cognition regimes. Because we endogenize the cognition order from choice (WOCI) and impose a cross-λ revealed-preference discipline (WARPD) that pins down how strict semiorder comparisons expand as cognition improves relative to a stable benchmark, we are able to make cross-λ statements meaningful without importing additional cardinal structure.

8.2 Revealed preference, WARP, and weakenings for set-valued choice

Beyond the semiorder tradition, our framework is related to a broader literature that explains violations of WARP through limits on deliberation, attention, and procedural constraints. A common theme is a two-stage interpretation of choice: the available menu is first reduced to a subset of "considered" alternatives, and only then is some criterion applied on that subset. In Manzini and Mariotti (2007)'s model of sequentially rationalizable choice, a fixed sequence of asymmetric "rationales" is applied to iteratively eliminate alternatives. In the limited-attention and limited-consideration tradition, the decision maker maximizes a stable preference only over an unobserved consideration set that can vary with the menu (Masatlioglu et al., 2012, Lleras et al., 2017). Cherepanov et al. (2013) propose a unifying "rationalization" framework in which a collection of psychological constraints determines which alternatives are rationalizable, and a fixed preference is then applied on that rationalizable set.

A key difference is that these filtering mechanisms can exclude an alternative that would otherwise be optimal under the agent's underlying ranking, simply because it is not shortlisted or not attended to. In our cognition-dependent representation, cognition does not act through a competing rationale or a menu-dependent filter. Instead, cognition determines which benchmark utility differences are behaviorally actionable, holding the benchmark ranking fixed. Consequently, benchmark maximizers remain admissible at every cognition level, while additional alternatives can enter the admissible set only because they lie within a cognition-dependent tolerance band around the benchmark best. Related work that relaxes the single-preference benchmark rather than constraining discrimination includes the multiple-rationales approach of Kalai et al. (2002), which measures the complexity of choice data by the minimal number of preference orderings needed for rationalization. Unlike our approach, these frameworks do not interpret across-environment variation as refinements of a single benchmark ranking.

8.2.1 Indifference, indecisiveness, and coarse indifference

Our interpretation of binary co-admissibility as a form of "coarse" choice is connected to the literature that links multi-valued choice to *incomplete* comparisons. In this strand, Eliaz and Ok (2006) introduce WARNI (Weak Axiom of Revealed Non-Inferiority) as a weakening of WARP that preserves the maximizing paradigm while allowing the rationalizing preference relation to be incomplete. On a finite-menu domain, they show that a choice correspondence satisfies WARNI if and only if it is rationalizable as undominated choice from a *unique regular* preorder, which need not be complete (see Theorem 2 and Corollary 1 in Eliaz and Ok, 2006). Regularity is a mild richness requirement that makes incomparability behaviorally meaningful and delivers uniqueness of the rationalizing preorder. In particular, WARNI provides a choice-theoretic foundation for distinguishing "true" indifference from indecisiveness (incompleteness) using menu variation beyond binary comparisons.

In our setting, *fixing a cognition level* λ, semiorder choice satisfies WARNI.[8] Thus each C_λ can be read, if one wishes, as undominated choice from an (in general incomplete) preorder in the sense of Eliaz and Ok (2006). Doing so refines our binary co-admissibility relation I_λ into two components: a transitive "indifference" relation $\sim_\lambda$ and a residual "incomparability" relation $\bowtie_\lambda$, where $I_\lambda = \sim_\lambda \dot\cup \bowtie_\lambda$. We record this refinement formally in section A3 of the online appendix, and show that $P_\lambda \cup \sim_\lambda$ is precisely the unique regular preorder singled out by Eliaz and Ok (2006). However, we do not emphasize this decomposition in the main text because, in our model, both $\sim_\lambda$ and $\bowtie_\lambda$ are resolution-induced objects: they shrink with cognition and are ultimately resolved at maximal cognition; the only remaining co-admissibilities coincide with genuine, transitive benchmark indifference (i.e. $I_{\bar\lambda} = \sim_{\bar\lambda}$, $\bowtie_{\bar\lambda} = \varnothing$).

Related models also connect incomplete comparisons to *deferral* or *status quo* behavior, rather than assuming all undominated options are admissible. For instance, Gerasimou (2018) studies choice deferral in the presence of incomplete (conflictual) comparisons, and Mu (2019) analyzes sequential procedures under incomplete preferences in which the decision maker may retain a status quo when comparisons are unresolved. These approaches use incompleteness to motivate *non-choice* or procedural stickiness, whereas we use cognition to discipline *graded admissibility* around a stable benchmark ranking.

8.3 Rational Inattention and information-based models

A complementary approach to limited discrimination keeps standard expected-utility preferences and explains apparent anomalies through an *unobserved* information-acquisition stage. In rational inattention, the decision maker optimally chooses what to attend to subject to an information-cost or capacity constraint (Sims, 2003). When the costs of obtaining information are modelled using the Shannon entropy, Matějka and McKay (2015) shows that

[8] On the finite-menu domain, the Jamison and Lau (1973) axioms JL1+JL3+JL5 (online appendix, section A2) imply semiorder-rationalizability Jamison and Lau (1973), Fishburn (1970b) and hence WARNI.

the choice probabilities take the familiar logit form. A large revealed-preference literature then asks which stochastic datasets can be rationalized by optimal costly information acquisition and how the underlying costs can be inferred from choice behavior.

Our framework shares with these papers the idea that "mistakes" arise from cognitive constraints rather than from taste shifts, but differs in where the constraint enters. In information models, mental processing operates through *state-dependent* learning and can induce genuine choice reversals across menus as attention optimally reallocates. By contrast, we study a *deterministic* comparative-statics: cognition indexes the resolution at which a fixed benchmark ranking is acted upon, so benchmark maximizers remain admissible at every cognition level. Additional options are admissible only when they are not too inferior compared to the benchmark preferences. In this sense, our model captures bounded discrimination as monotone refinement of resolved comparisons, rather than as endogenous information acquisition that can reorganize behavior across menus.

8.4 Fuzzy preference relations, vagueness, and the Sorites paradox

Fuzzy set theory was introduced to model vagueness via graded membership (Zadeh, 1965). Fuzzy binary relations extend this idea to pairwise comparisons and have a long tradition in decision analysis and preference aggregation, often imposing transitivity-like restrictions defined through t-norms and related operators (Orlovsky, 1978, Dubois and Prade, 1980, Ovchinnikov, 1981, Fodor and Roubens, 1994, Klement et al., 2000). Our use of fuzziness differs in interpretation: unlike previous work (where fuzzy relations are primitives), the fuzzy relation $R(x, y)$ is a revealed-preference object encoding the (ordinal) difficulty of recognizing the benchmark strict statement "$y \succ x$." Cognition levels then correspond to calibrated cuts of R (Definition 12), so that higher cognition reveals more strict comparisons and yields smaller admissible sets.

A conceptual contribution is to show that two canonical responses to the Sorites paradox coincide once expressed in this revealed-preference language. On one hand, fuzzy-logic approaches treat vague predicates as holding to a degree rather than being simply true or false (Zadeh, 1965, Keefe, 2000). On the other hand, JND and semiorder approaches deny that local indistinguishability is transitive, so chains of tolerable differences do not force endpoint identity (Luce, 1956, Fishburn, 1985). Section 5.5 makes precise that under fuzzy rationalizability these are two descriptions of the same cut structure. This perspective also interfaces naturally with broader philosophical discussions of vagueness (including supervaluationist and epistemic views) (Fine, 1975, Williamson, 1994), though our aim is not to adjudicate among them but to isolate the operational content relevant for revealed preference.

9 Conclusion

This paper studies choice when a decision-maker has a stable benchmark ranking of alternatives, but the ability to act on small differences varies across cognition regimes. The primitive is a family of admissibility correspondences $\mathcal{C} = \{C_\lambda : \lambda \in \Lambda\}$, where λ indexes cognitive resolution. We begin with the cognition-dependent representation of Definition 1, in which a single benchmark utility u rationalizes the entire family and cognition operates through tolerances $\varepsilon(\lambda, x)$ that weakly shrink with cognition and vanish at maximal cognition. We then establish equivalences between CDR and (i) the choice axioms WOCI and WARPD, (ii) a preference-side characterization in terms of nested semiorders with ordered discovery, and (iii) a fuzzy-preferences formulation in which cognition levels correspond to cuts of a fuzzy difficulty relation.

Two lessons may be useful for applications. First, the framework separates tastes from cognition: the benchmark ranking is stable, and cognition affects only discriminability. This supports comparative statics across environments (time pressure, noise, expertise, incentives) without interpreting these environments as taste shifters. Second, the equivalence between semiorder, threshold, and fuzzy formulations suggests different empirical entry points: depending on what is observed or elicited, one can work with a nested family of semiorders, estimate a cognition-dependent tolerance profile, or directly recover a fuzzy difficulty relation whose calibrated cuts reproduce admissibility behavior.

Several directions for future work are natural. One is to incorporate stochastic choice explicitly, interpreting $C_\lambda(S)$ as the support (or a high-probability set) of a random choice rule and studying how the axioms interact with random-utility and consideration-based models (Block and Marschak, 1960, Luce, 1959, McFadden, 1974, Manzini and Mariotti, 2014). A second is to endogenize cognition by linking the calibration $\nu(\lambda)$ to optimal attention or information acquisition (Sims, 2003, Caplin and Dean, 2015, Matějka and McKay, 2015, Masatlioglu et al., 2012). A third is to study welfare and policy questions under the maintained separation between tastes and discrimination: interventions that change cognition regimes (time, incentives, noise) need not change benchmark preferences, but can change which alternatives are admissible and therefore selected.

References

Balart, P. (2021): "Semiorder preferences and price-oriented buyers in a Hotelling mode," *Journal of Economic Behavior & Organization*, 188, 394–407.

Beja, A. and I. Gilboa (1992): "Numerical Representations of Imperfectly Ordered Preferences (a Unified Geometric Exposition)," *Journal of Mathematical Psychology*, 36, 426–449.

Block, H. D. and J. Marschak (1960): "Random Orderings and Stochastic Theories of Responses," in *Contributions to Probability and Statistics: Essays in Honor of Harold Hotelling*, ed. by I. Olkin, S. G. Ghurye, W. Hoeffding, W. G. Madow, and H. B. Mann, Stanford, CA: Stanford University Press, 97–132.

Caplin, A. and M. Dean (2015): "Revealed Preference, Rational Inattention, and Costly Information Acquisition," *American Economic Review*, 105, 2183–2203.

Cherepanov, V., T. Feddersen, and A. Sandroni (2013): "Rationalization," *Theoretical Economics*, 8, 775–800.

Dubois, D. and H. Prade (1980): *Fuzzy Sets and Systems: Theory and Applications*, New York: Academic Press.

Dziewulski, P. (2020): "Just-noticeable Difference as a Behavioural Foundation of the Critical Cost-efficiency Index," *Journal of Economic Theory*, 188, 105071.

Eliaz, K. and E. A. Ok (2006): "Indifference or Indecisiveness? Choice-Theoretic Foundations of Incomplete Preferences," *Games and Economic Behavior*, 56, 61–92.

Fechner, G. T. (1860): *Elemente der Psychophysik*, Leipzig: Breitkopf und Härtel.

Fine, K. (1975): "Vagueness, Truth and Logic," *Synthese*, 30, 265–300.

Fishburn, P. C. (1970a): ""Intransitive Indifference with Unequal Indifference Intervals"," *Journal of Mathematical Psychology*, 7, 144–149.

——— (1970b): *Utility Theory for Decision Making*, New York: Wiley.

——— (1985): *Interval Orders and Interval Graphs*, New York: Wiley.

Fodor, J. and M. Roubens (1994): *Fuzzy Preference Modelling and Multicriteria Decision Support*, Dordrecht: Kluwer Academic Publishers.

Gerasimou, G. (2018): "Indecisiveness, Undesirability and Overload Revealed Through Rational Choice Deferral," *The Economic Journal*, 128, 2450–2479.

Jamison, D. T. and L. J. Lau (1973): "Semiorders and the Theory of Choice," *Econometrica*, 41, 901–912.

Kalai, G., A. Rubinstein, and R. Spiegler (2002): "Rationalizing Choice Functions by Multiple Rationales," *Econometrica*, 70, 2481–2488.

Keefe, R. (2000): *Theories of Vagueness*, Cambridge: Cambridge University Press.

Klement, E. P., R. Mesiar, and E. Pap (2000): *Triangular Norms*, Dordrecht: Kluwer Academic Publishers.

Liu, Y. and E. E. Kerre (1998): "An Overview of Fuzzy Quantifiers. (I). Interpretations," *Fuzzy Sets and Systems*, 95, 1–21.

Lleras, J. S., Y. Masatlioglu, D. Nakajima, and E. Y. Ozbay (2017): "When More Is Less: Limited Consideration," *Journal of Economic Theory*, 170, 70–85.

Luce, R. D. (1956): "Semiorders and a Theory of Utility Discrimination," *Econometrica*, 24, 178–191.

——— (1959): *Individual Choice Behavior: A Theoretical Analysis*, New York: Wiley.

Manzini, P. and M. Mariotti (2007): "Sequentially Rationalizable Choice," *American Economic*

Review, 97, 1824–1839.

——— (2012): "Choice by Lexicographic Semiorders," *Journal of Economic Theory*, 147, 2196–2220.

——— (2014): "Stochastic Choice and Consideration Sets," *Econometrica*, 82, 1153–1176.

Masatlioglu, Y., D. Nakajima, and E. Y. Ozbay (2012): "Revealed Attention," *American Economic Review*, 102, 2183–2205.

Matějka, F. and A. McKay (2015): "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model," *American Economic Review*, 105, 272–298.

McFadden, D. (1974): "Conditional Logit Analysis of Qualitative Choice Behavior," in *Frontiers in Econometrics*, ed. by P. Zarembka, New York: Academic Press, 105–142.

Mu, X. (2019): "Amendment Voting with Incomplete Preferences," Working Papers 2019-29, Princeton University. Economics Department.

Nishimura, H. (2018): "The Transitive Core: Inference of Welfare from Nontransitive Preference Relations," *Theoretical Economics*, 13, 579–606.

Orlovsky, S. A. (1978): "Decision-making with a Fuzzy Preference Relation," *Fuzzy Sets and Systems*, 1, 155–167.

Ovchinnikov, S. V. (1981): "Structure of Fuzzy Binary Relations," *Fuzzy Sets and Systems*, 6, 169–195.

Roberts, F. S. (1971): "Homogeneous families of semiorders and the theory of probabilistic consistency," *Journal of Mathematical Psychology*, 8, 248–263.

Scott, D. and P. Suppes (1958): "Foundational Aspects of Theories of Measurement," *The Journal of Symbolic Logic*, 23, 113–128.

Shimoda, M. (2002): "A Natural Interpretation of Fuzzy Sets and Fuzzy Relations," *Fuzzy Sets and Systems*, 128, 135–147.

Sims, C. A. (2003): "Implications of Rational Inattention," *Journal of Monetary Economics*, 50, 665–690.

Weber, E. H. (1834): *De Pulsu, Resorptione, Auditu et Tactu*.

Williamson, T. (1994): *Vagueness*, London: Routledge.

Zadeh, L. A. (1965): "Fuzzy Sets," *Information and Control*, 8, 338–353.

——— (1971): "Similarity Relations and Fuzzy Orderings," *Information Sciences*, 3, 177–200.

A Appendix: Proofs

A.1 Proofs and auxiliary results: characterization and main theorem

Proposition 6 (Ordering Λ)**.** *If the family $\{C_\lambda : \lambda \in \Lambda\}$ satisfies WOCI, then the relation $\geq$ is a weak order on Λ.*

Proof. The relation $\geq$ is a weak order if it is complete and transitive.

Completeness. Fix $\lambda, \lambda' \in \Lambda$. Suppose that $\lambda' \not\geq \lambda$. By definition, this means there exists some $S \in \mathcal{S}$ such that $C_{\lambda'}(S) \not\subseteq C_\lambda(S)$. By Nestedness (WOCI(i)), it must be that $C_\lambda(S) \subsetneq C_{\lambda'}(S)$. Then, by Directional Consistency (WOCI(ii)), it follows that $C_\lambda(T) \subseteq C_{\lambda'}(T)$ for all $T \in \mathcal{S}$.

Transitivity. Suppose $\lambda'' \geq \lambda'$ and $\lambda' \geq \lambda$. By definition, this means $C_{\lambda''}(S) \subseteq C_{\lambda'}(S)$ for all $S \in \mathcal{S}$, and $C_{\lambda'}(S) \subseteq C_\lambda(S)$ for all $S \in \mathcal{S}$. By the transitivity of set inclusion, this implies $C_{\lambda''}(S) \subseteq C_\lambda(S)$ for all $S \in \mathcal{S}$. ∎

Proposition 7. *If the family $\{C_\lambda : \lambda \in \Lambda\}$ satisfies WOCI and X is finite, then there exists a finite number of $\geq$-equivalence classes on Λ.*

Proof. Since X is finite, the set of menus $\mathcal{S} = \{S_1, \dots, S_n\}$ is also finite. For each $S_i \in \mathcal{S}$, we can define a function $g_i : \Lambda \to \mathcal{P}(S_i)$ by $g_i(\lambda) = C_\lambda(S_i)$. Since S_i is a finite set, its power set $\mathcal{P}(S_i)$ is finite. Therefore, the image of g_i, denoted $\operatorname{Im} g_i$, must also be finite.

Let $\{C_i^1, C_i^2, \dots, C_i^L\}$ be the finite set of possible choice sets for S_i (i.e., the image of g_i). The inverse images $\Lambda_i^\ell := g_i^{-1}(C_i^\ell)$ for $\ell = 1, \dots, L_i$ form a finite partition of Λ.

We can now construct a new partition, $\hat{\Lambda}$, which is the common refinement of the partitions for all menus $S_1, \dots, S_n$:

$$\hat{\Lambda} = \left\{ \bigcap_{i=1}^{n} \Lambda_i^{\ell_i} : \ell_i \in \{1, \dots, L_i\} \text{ for each } i \in \{1, \dots, n\} \right\} \setminus \{\varnothing\}.$$

This partition $\hat{\Lambda}$ is also finite. By construction, if any two indices λ and λ' belong to the same cell in $\hat{\Lambda}$, then $C_\lambda(S_i) = C_{\lambda'}(S_i)$ for all $i = 1, \dots, n$. This implies that $C_\lambda(S) = C_{\lambda'}(S)$ for all $S \in \mathcal{S}$.

Two indices belong to the same $\geq$-equivalence class if and only if $\lambda \sim \lambda'$, which is equivalent to $C_\lambda(S) = C_{\lambda'}(S)$ for all $S \in \mathcal{S}$. Therefore, the cells of the finite partition $\hat{\Lambda}$ are the $\geq$-equivalence classes, and there must be a finite number of them. ∎

Theorem 1. *For a family of choice correspondences $\mathcal{C} = \{C_\lambda : \lambda \in \Lambda\}$, the following are equivalent:*

(i) $\mathcal{C}$ admits a cognition-dependent representation (Definition 1).

(ii) $\mathcal{C}$ satisfies WOCI and WARPD.

Proof. WARPD **and** WOCI $\Longrightarrow$ **cognition-dependent representation.**

Assume the family of choice correspondences $\mathcal{C}$ satisfies both WARPD and WOCI. We construct the functions u and ε and show they satisfy the required properties.

By WOCI, the weak order $\geq$ on Λ allows us to identify a maximal cognition index, $\overline{\lambda}$. At this level, WARPD implies that the choice correspondence $C_{\overline{\lambda}}$ satisfies the standard WARP. Because X is finite, this guarantees the existence of a utility function $u : X \to \mathbb{R}$ that represents $C_{\overline{\lambda}}$, such that $C_{\overline{\lambda}}(S) = \{x \in S : u(x) \geq u(y) \text{ for all } y \in S\}$. This u will serve as our rational benchmark utility.

Next, for any $\lambda \in \Lambda$ and $x \in X$, define the threshold function ε as follows:

$$\varepsilon(\lambda, x) := \max_{z \in X}\{u(z) - u(x) \ : \ x \in C_\lambda(\{x, z\})\}.$$

This function captures the maximum utility difference that alternative x can tolerate from a competitor z and still be chosen in a binary contest at cognition level λ. Since $x \in C_\lambda(\{x, x\})$, the set is non-empty and well-defined. By construction, $\varepsilon(\lambda, x) \geq u(x) - u(x) = 0$.

We must show that for any $S \in \mathcal{S}$, $x \in C_\lambda(S) \iff u(y) - u(x) \leq \varepsilon(\lambda, x)$ for all $y \in S$. We prove each direction of the equivalence.

($\Longleftarrow$) Assume for all $y \in S$, $u(y) - u(x) \leq \varepsilon(\lambda, x)$. We must show $x \in C_\lambda(S)$.

The proof is by contradiction. Assume $x \notin C_\lambda(S)$. Suppose $z \in \bigcap_{\lambda'} C_{\lambda'}(S) = C_{\overline{\lambda}}(S)$ (by WOCI). Since WARPD implies CRβ (as shown in Proposition 2 in the online appendix, it must be that $x \notin C_\lambda(\{x, z\})$.

By definition of ε, if $x \notin C_\lambda(\{x, z\})$, it must be that $u(z) - u(x) > \varepsilon(\lambda, x)$. Otherwise, there would exist a z' such that $u(z') > u(z)$ and $x \in C_\lambda(\{x, z'\})$. That would imply $x \in C_\lambda(\{x, z, z'\})$ and thus that $x \in C_\lambda(\{x, z\})$. That is impossible, so it must be that, indeed, $u(z) - u(x) > \varepsilon(\lambda, x)$. But this contradicts our initial assumption that for all $y \in S$, $u(y) - u(x) \leq \varepsilon(\lambda, x)$. Thus, the original assumption must be false, and $x \in C_\lambda(S)$.

($\Longrightarrow$) Assume $x \in C_\lambda(S)$. We must show that, for all $y \in S$, $u(y) - u(x) \leq \varepsilon(\lambda, x)$.

The proof is again by contradiction. Assume there exists a $y \in S$ such that $u(y) - u(x) > \varepsilon(\lambda, x)$. By definition of ε, this immediately implies that $x \notin C_\lambda(\{x, y\})$.

However, we assumed $x \in C_\lambda(S)$. Since $\{x, y\} \subseteq S$, and since WARPD implies Uniform Property α, it must be that $x \in C_\lambda(\{x, y\})$. This is a contradiction. Thus, our assumption must be false, and for all $y \in S$, we must have $u(y) - u(x) \leq \varepsilon(\lambda, x)$.

Next, we show that the remaining properties of the cognition-dependent representation hold.

(i) $\varepsilon(\overline{\lambda}, x) = 0$ **for all** $x \in X$**.** At $\lambda = \overline{\lambda}$, $C_{\overline{\lambda}}$ is represented by u. Thus, $x \in C_{\overline{\lambda}}(\{x, z\})$ if and only if $u(x) \geq u(z)$, which means $u(z) - u(x) \leq 0$. The maximum such value is 0.

(ii) $\varepsilon(\cdot, x)$ **is nonincreasing for all** $x \in X$**.** Fix $\lambda < \lambda'$. By WOCI there exists $S \in \mathcal{S}$ and an $x \in C_\lambda(S)$ such that $x \notin C_{\lambda'}(S)$. Thus, there exists $z \in S$ such that $u(z) - u(x) \leq \varepsilon(\lambda, x)$ but $u(z) - u(x) > \varepsilon(\lambda', x)$. Together, these inequalities imply that $\varepsilon(\lambda', x) < \varepsilon(\lambda, x)$.

(iii) **for all** $x, y \in X$**,**$\lambda \in \Lambda$**,** $u(x) \geq u(y)$**, then** $u(x) + \varepsilon(\lambda, x) \geq u(y) + \varepsilon(\lambda, y)$**.** Suppose, by way of contradiction, that $u(x) \geq u(y)$ but $u(x) + \varepsilon(\lambda, x) < u(y) + \varepsilon(\lambda, y)$ for some x, $y \in X$ and $\lambda \in \Lambda$. Then, there exists z_y such that $u(z_y) > u(x) \geq u(y)$, $y \in C_\lambda(\{y, z_y\})$ but $x \notin C_\lambda(\{x, z_y\})$. Since $u(z_y) > u(x) \geq u(y)$, $z_y \in C_{\overline{\lambda}}(\{x, y, z_y\})$. By WARPD, it must then be that $C_\lambda(\{x, y, z_y\}) = \{y, z_y\}$. But because $x \in C_{\overline{\lambda}}(\{x, y\})$ and $y \in C_\lambda(\{x, y, z_y\})$, WARPD also implies that x is in $C_\lambda(\{x, y, z_y\})$, a contradiction.

Cognition-dependent representation $\Longrightarrow$ **WARPD and WOCI.**

Suppose the family $\mathcal{C} := \{C_\lambda : \lambda \in \Lambda\}$ admits a cognition-dependent representation. We show that WOCI and WARPD hold.

- WOCI holds.

  Nestedness. Let $\lambda, \lambda' \in \Lambda$. Assume $\lambda' > \lambda$. By property (ii) of the representation, $\varepsilon(\lambda', x) \leq \varepsilon(\lambda, x)$ for all x. Now, if $x \in C_{\lambda'}(S)$, then for all $y \in S$, $u(y) - u(x) \leq \varepsilon(\lambda', x)$. Since $\varepsilon(\lambda', x) \leq \varepsilon(\lambda, x)$, it is also true that $u(y) - u(x) \leq \varepsilon(\lambda, x)$, which implies $x \in C_\lambda(S)$. Thus, $C_{\lambda'}(S) \subseteq C_\lambda(S)$.

  Consistency. If $C_\lambda(T) \subsetneq C_{\lambda'}(T)$ for some T, that means that there exists $x, y \in T$ such that $\varepsilon(\lambda', x) > u(y) - u(x) > \varepsilon(\lambda, x)$, so that $\lambda \geq \lambda'$. By Nestedness, $C_\lambda(S) \subseteq C_{\lambda'}(S)$ for all S.

- WARPD holds. Let $\lambda \in \Lambda$, Since WOCI holds, the intersection condition $\bigcap C_{\lambda'}$ can be replaced by $C_{\overline{\lambda}}$. We assume $x \in C_{\overline{\lambda}}(S)$ and $y \in C_\lambda(T)$ for $\{x, y\} \subseteq S \cap T$.

  From $x \in C_{\overline{\lambda}}(S)$ and property (i) of the representation, we know $\varepsilon(\overline{\lambda}, x) = 0$. Thus, for all $z \in S$, $u(z) - u(x) \leq 0$, which means $u(x) = \max_{z \in S} u(z)$. From $y \in C_\lambda(T)$, we know that for all $z \in T$, $u(z) - u(y) \leq \varepsilon(\lambda, y)$, which is equivalent to $u(y) + \varepsilon(\lambda, y) \geq \max_{z \in T} u(z)$.

  $x \in C_\lambda(T)$**.** Since $y \in S$, we have $u(x) \geq u(y)$. By property (iii) (Regularity), this implies $u(x) + \varepsilon(\lambda, x) \geq u(y) + \varepsilon(\lambda, y)$. Combining this with what we know about y, we get $u(x) + \varepsilon(\lambda, x) \geq u(y) + \varepsilon(\lambda, y) \geq \max_{z \in T} u(z)$. This proves $x \in C_\lambda(T)$.

  $y \in C_\lambda(S)$**.** Since $x \in T$, we have $u(y) + \varepsilon(\lambda, y) \geq u(x)$. We also know $u(x) \geq \max_{z \in S} u(z)$. Combining these gives $u(y) + \varepsilon(\lambda, y) \geq \max_{z \in S} u(z)$. This proves $y \in C_\lambda(S)$.

■

A.2 Proofs and auxiliary results: Preferences

Lemma 2. *Let X be finite and let $\mathcal{P} = \{P_\lambda : \lambda \in \Lambda\}$ satisfy N. Then there exists a relation $P_{\overline{\lambda}}$ such that $P_\lambda \subseteq P_{\overline{\lambda}}$ for all λ. Moreover, without loss of generality we may take $\overline{\lambda} \in \Lambda$ by replacing Λ with the finite set of distinct relations $\{P_\lambda : \lambda \in \Lambda\}$ ordered by inclusion.*

Proof. Since X is finite, there are only finitely many distinct binary relations on X. Nestedness N implies that the set $\{P_\lambda : \lambda \in \Lambda\}$ is a chain under inclusion, hence has a maximal element $P_{\overline{\lambda}}$ among its finitely many distinct members. Relabel the distinct relations by a (finite) chain of indices to ensure $\overline{\lambda} \in \Lambda$. ∎

Lemma 3. *Assume GA. Then the relation $\succeq$ defined in (8) is complete and reflexive. If in addition N holds and $\overline{\lambda}$ is maximal, then*

$$x \succeq y \iff y \not P_{\overline{\lambda}} x.$$

Proof. Reflexivity is immediate since each P_λ is asymmetric, so $\neg(xP_\lambda x)$ for all λ. For completeness, fix distinct x, y. If neither $x \succeq y$ nor $y \succeq x$ held, then there would exist λ, μ with $yP_\lambda x$ and $xP_\mu y$, contradicting GA. For the final claim, use Lemma 2: under nestedness, $yP_\lambda x$ for some λ is equivalent to $yP_{\overline{\lambda}} x$. ∎

Theorem 2. *For a family of choice correspondences $\mathcal{C} = \{C_\lambda : \lambda \in \Lambda\}$, the following are equivalent:*

(i) $\mathcal{C}$ admits a cognition-dependent representation in the sense of Definition 1.

(ii) $\mathcal{C}$ is $\mathcal{P}$–rationalizable (Definition 4) by a family $\mathcal{P} = \{P_\lambda : \lambda \in \Lambda\}$ satisfying GA, N, UW, and one (hence all) of OD, SSC, WSSC.

(iii) $\mathcal{C}$ is $\mathcal{P}$–rationalizable (Definition 4) by a homogeneous family of semiorders (Definition 7) $\mathcal{P} = \{P_\lambda : \lambda \in \Lambda\}$ satisfying GA, N, UW.

For proving the Theorem, we will use the following Lemma:

Lemma 4. *Let X be finite and let $\mathcal{P} = \{P_\lambda : \lambda \in \Lambda\}$ be nested by inclusion with maximal element $P_{\overline{\lambda}}$. Suppose $\mathcal{P}$ satisfies GA,N and UW, plus one of the conditions of the Theorem (OD,SSC,WSSC or Definition 7). Let $\succeq$ be the weak order induced by $P_{\overline{\lambda}}$: $x \succeq y \iff y \not P_{\overline{\lambda}} x$, and let $\succ$ denote its strict part.*[9]

Assume that for a fixed λ, P_λ is compatible with $\succeq$ and satisfies:

$$xP_\lambda y \implies x \succ y. \tag{14}$$

Let u be any numerical representation of $\succeq$. For each $x \in X$, define

$$U_\lambda(x) := \max\{u(y) : \neg(yP_\lambda x)\}, \qquad \varepsilon(\lambda, x) := U_\lambda(x) - u(x) \geq 0.$$

[9] We know from Roberts (1971) that $\succeq$ is transitive as well as complete, and is thus a weak order.

Then for all $x, y \in X$,

$$y \not P_\lambda x \iff u(y) - u(x) \le \varepsilon(\lambda, x).$$

Moreover, $u(x) \ge u(y)$ *implies* $U_\lambda(x) \ge U_\lambda(y)$ *or, equivalently* $u(x) + \varepsilon(\lambda, x) \ge u(y) + \varepsilon(\lambda, y)$.

Proof of Lemma. Fix λ.

Nonemptiness holds since P_λ is asymmetric, so $\neg(xP_\lambda x)$. To see it is a lower set, it is equivalent to show that the strict upper contour $\{y : yP_\lambda x\}$ is an upper set. So take $yP_\lambda x$ and $z \ge y$; we prove $zP_\lambda x$. By (14), $y > x$, hence also $z > x$ (since $\ge$ is transitive). If $\neg(zP_\lambda x)$, then either $xP_\lambda z$ or $zI_\lambda x$. The first is impossible because $xP_\lambda z$ would imply $x > z$, contradicting $z > x$. If instead $zI_\lambda x$, then $z \ge y \ge x$ (because $yP_\lambda x \Rightarrow y \ge x$ by compatibility) and compatibility (ii) yields $yI_\lambda x$, contradicting $yP_\lambda x$. Thus $zP_\lambda x$, and the upper contour is an upper set.

Since X is finite and the nonwinner set is nonempty, $U_\lambda(x)$ is well-defined. If $\neg(yP_\lambda x)$, then by definition $u(y) \le U_\lambda(x)$, hence $u(y) - u(x) \le \varepsilon(\lambda, x)$. Conversely, suppose $u(y) - u(x) \le \varepsilon(\lambda, x)$, i.e. $u(y) \le U_\lambda(x)$. Pick $y^* \in X$ such that $\neg(y^* P_\lambda x)$ and $u(y^*) = U_\lambda(x)$. Then $u(y) \le u(y^*)$ implies $y \le y^*$ in $\ge$. Since the nonwinner set is a lower set (Step 1) and y^* is a nonwinner, y is also a nonwinner: $\neg(yP_\lambda x)$.

Let $u(x) \ge u(y)$, i.e. $x \ge y$. We show $\{z : \neg(zP_\lambda x)\} \supseteq \{z : \neg(zP_\lambda y)\}$. Equivalently, take $zP_\lambda x$ and prove $zP_\lambda y$. Since $zP_\lambda x$, (14) gives $z > x$; with $x \ge y$ this implies $z > y$. If $\neg(zP_\lambda y)$ then either $yP_\lambda z$ (impossible since it would give $y > z$) or $zI_\lambda y$. But $z \ge x \ge y$ and $zI_\lambda y$ would imply $zI_\lambda x$ by compatibility (ii), contradicting $zP_\lambda x$. Hence $zP_\lambda y$. Therefore the nonwinner set for x contains that for y, and taking maxima yields $U_\lambda(x) \ge U_\lambda(y)$. ∎

Proof of Theorem 2. (i)⇒(iii). Assume $\mathcal{C}$ admits a CDR: there exist u and $\varepsilon(\lambda, \cdot)$ as in Definition 1. Define $\ge$ by $x \ge y \iff u(x) \ge u(y)$. For each λ, define P_λ by binary admissibility:

$$xP_\lambda y \quad\iff\quad y \notin C_\lambda(\{x, y\}) \quad\iff\quad u(x) > u(y) + \varepsilon(\lambda, y).$$

Then $C_\lambda = \max_{P_\lambda}$, so $\mathcal{C}$ is $\mathcal{P}$–rationalizable.

Nestedness of tolerances in Definition 1 implies nestedness of strict comparisons: if $xP_\lambda y$ and $\lambda' \ge \lambda$ (higher cognition), then $u(x) > u(y) + \varepsilon(\lambda, y) \ge u(y) + \varepsilon(\lambda', y)$, hence $xP_{\lambda'} y$. Thus N holds (after ordering indices by inclusion as usual). General asymmetry GA holds because $xP_\lambda y$ implies $u(x) > u(y)$, hence $u(y) \not> u(x) + \varepsilon(\mu, x)$ for any μ.

Uniform witnessing UW follows from the regularity built into the CDR: if $xI_\lambda y$ for all λ, then in particular $xI_{\overline{\lambda}} y$ at maximal cognition, so $u(x) = u(y)$. Regularity then implies $\varepsilon(\lambda, x) = \varepsilon(\lambda, y)$ for each λ (because $u(x) = u(y)$ forces $u(x) + \varepsilon(\lambda, x) = u(y) + \varepsilon(\lambda, y)$), hence all binary comparisons against any z coincide.

Finally, each P_λ is compatible with $\ge$: compatibility (i) holds since $xP_\lambda y \Rightarrow u(x) > u(y)$. For (ii), take $x \ge y \ge z$ and $xI_\lambda z$. Then $xI_\lambda z$ means both $\neg(xP_\lambda z)$ and $\neg(zP_\lambda x)$, i.e. $u(x) \le u(z) + \varepsilon(\lambda, z)$ and $u(z) \le u(x) + \varepsilon(\lambda, x)$. Since $u(x) \ge u(y) \ge u(z)$ and $u(y) +$

$\varepsilon(\lambda, y) \geq u(z) + \varepsilon(\lambda, z)$ by regularity, we get $u(x) \leq u(y) + \varepsilon(\lambda, y)$; the other inequality $u(y) \leq u(x) + \varepsilon(\lambda, x)$ is immediate from $u(y) \leq u(x)$. Thus $xI_\lambda y$. A symmetric argument gives $yI_\lambda z$. Hence P_λ is compatible with $\succeq$ for every λ. This is exactly a homogeneous family satisfying the desired properties.

(iii)⇒(ii). If $\mathcal{P}$ is a homogeneous nested family of semiorders, then N holds by definition. Moreover, since the benchmark $\succeq$ is induced by the maximal relation, strict comparisons cannot reverse across regimes, so GA holds. Uniform witnessing UW is a standard implication of compatibility with a weak order when $xI_\lambda y$ holds at all regimes (and hence $x \sim y$ in the benchmark). Finally, by Proposition 2, under GA–UW homogeneity is equivalent to each of OD, SSC, WSSC. Thus (ii) holds.

(ii)⇒(i). Assume (ii). By Proposition 2, the benchmark relation $\succeq$ defined by (8) is a weak order and is compatible with each P_λ. Let u be any numerical representation of $\succeq$. Let $\overline{\lambda}$ be maximal (Lemma 2); then $P_{\overline{\lambda}}$ is the strict part of $\succeq$.

For each λ and each x, define $U_\lambda(x)$ and $\varepsilon(\lambda, x)$ as in Lemma 4. Condition (14) holds because nesting implies $xP_\lambda y \Rightarrow xP_{\overline{\lambda}} y$, hence $x \succ y$ in the benchmark. Therefore Lemma 4 applies and yields:

$$\neg(yP_\lambda x) \iff u(y) - u(x) \leq \varepsilon(\lambda, x).$$

Using $\mathcal{P}$–rationalizability, for any menu S,

$$x \in C_\lambda(S) \iff x \in \max_{P_\lambda}(S) \iff (\forall y \in S)\neg(yP_\lambda x) \iff (\forall y \in S)u(y) - u(x) \leq \varepsilon(\lambda, x),$$

which is the threshold admissibility form of Definition 1.

Nesting N implies $\varepsilon(\lambda, x)$ weakly shrinks with cognition: if $P_\lambda \subseteq P_{\lambda'}$, then $\{y : \neg(yP_{\lambda'}x)\} \subseteq \{y : \neg(yP_\lambda x)\}$, hence $U_{\lambda'}(x) \leq U_\lambda(x)$ and $\varepsilon(\lambda', x) \leq \varepsilon(\lambda, x)$. At maximal cognition, since $P_{\overline{\lambda}}$ is the strict part of $\succeq$, $\neg(yP_{\overline{\lambda}}x)$ holds if and only if $u(y) \leq u(x)$, so $U_{\overline{\lambda}}(x) = u(x)$ and $\varepsilon(\overline{\lambda}, x) = 0$. Finally, Lemma 4 also gives the required regularity: $u(x) \geq u(y) \Rightarrow u(x) + \varepsilon(\lambda, x) \geq u(y) + \varepsilon(\lambda, y)$. Thus (u, ε) is a CDR. ∎

A.3 Proofs and auxiliary results: Fuzzy preferences

Theorem 3. *For a family of choice correspondences $\mathcal{C} = \{C_\lambda : \lambda \in \Lambda\}$, the following are equivalent:*

(i) $\mathcal{C}$ admits a cognition-dependent representation (Definition 1).

(ii) $\mathcal{C}$ is fuzzy regular-rationalizable. (Definition 12).

Proof. We will prove it using Theorem 1. That is, we will show that our rationalizability statement is equivalent to WARPD and WOCI.

(i) $\implies$ (ii)

Suppose WARPD and WOCI hold, and take Λ to be weakly ordered with maximal element $\overline{\lambda}$. Let $\nu : \Lambda \implies (0, 1]$ be any calibration with $\nu(\overline{\lambda}) = \bar{v}$.

Let $R(x,y) = \max\{\nu(\lambda) : x \in C_\lambda(\{x,y\})\}$.

We must show that R is fuzzy rational and that, for all S, $C_\lambda(S) = C^{R;\nu(\lambda)}(S)$.

Because each C_λ is nonempty, $C_{\overline{\lambda}}(\{x,y\})$ contains either x, or y, or both. Then, $R(x,y) = \bar{v}$ or $R(y,x) = \bar{v}$, i.e. R is fuzzy complete.

Now suppose $R(x,y) = \bar{v}$ and $R(y,z) = \bar{v}$ and consider $\lambda \in \max\{\mu \in \Lambda : z \in C_\mu(\{x,z\})\}$. Note that $R(z,x) = \nu(\lambda)$, since ν is strictly increasing. Successive applications of WARPD yield: $z \in C_\lambda(\{x,y,z\}) \implies z \in C_\lambda(\{y,z\}) \implies y \in C_\lambda(\{x,y,z\}) \implies y \in C_\lambda(\{x,y\})$. That means that $R(y,x) \geq \alpha(\lambda)$ and $R(z,y) \geq \alpha(\lambda)$ both hold, and thus R is fuzzy transitive.

Finally, we show that $C_\lambda(S) = C^{R;\nu(\lambda)}(S)$ for every $S \in \mathcal{S}$. Let $x \in C_\lambda(S)$. By WARPD, $x \in C_\lambda(\{x,y\})$ for all $y \in S$. Thus, $R(x,y) \geq \nu(\lambda)$ for all $y \in S$. But then, $x \in C^{R;\nu(\lambda)}(S)$.

Conversely, suppose $x \in C^{R;\nu(\lambda)}(S)$. Then, $R(x,y) \geq \nu(\lambda)$ for all $y \in S$. That means $x \in C_\lambda(\{x,y\})$ for all $y \in S$. By WARPD, we have that $x \in C_\lambda(S)$.

(ii) $\implies$ **(i)** Let Λ be ordered according to ν. That is, say that $\lambda' > \lambda$ whenever $\nu(\lambda') > \nu(\lambda)$. WOCI clearly holds since the α-cuts are nested. Moreover, $\nu(\overline{\lambda}) = \bar{v}$ for some $\overline{\lambda}$; since $1 \geq \alpha(\lambda)$ for all λ, $\overline{\lambda} \geq \lambda$ for all λ.

For WARPD, fix $x, y \in S \cap T$ with $x \in C_{\overline{\lambda}}(S)$ and $y \in C_\lambda(T)$.

Suppose $x \notin C_\lambda(T)$. Then, there exists $z_T \in T$ such that $R(x,z_T) < \nu(\lambda)$. By fuzzy completeness, $R(z_T,x) = \bar{v}$. Since $x \in C_{\overline{\lambda}}(S)$, $R(x,y) \geq \nu(\overline{\lambda}) = \bar{v} \implies R(x,y) = 1$.

Then, by fuzzy transitivity, $R(y,z_T) \leq R(x,z_T) < \nu(\lambda)$; thus $y \notin nC_\lambda(T)$, a contradiction. Thus, $x \in C_\lambda(T)$.

Now suppose $y \notin C_\lambda(S)$. That is, there exists $z_S \in S$ such that $R(y,z_S) < \nu(\lambda)$. That means $R(z_S,y) = \bar{v}$, by fuzzy completeness. Moreover, $R(x,z_S) = \bar{v}$, since $x \in C_{\overline{\lambda}}(S)$. By fuzzy transitivity, we have that $R(y,x) \leq R(y,z_S) < \nu(\lambda)$. But then y could not be in $C_\lambda(T)$, a contradiction. ∎

A.4 Proofs and auxiliary results: Incompleteness of observation

Proposition 4. *Let $\tilde{\Lambda} \subseteq \Lambda$ be nonempty. Define a binary relation $T_{\tilde{\Lambda}}$ on X by, for $x, y \in X$,*

$$
\begin{aligned}
xT_{\tilde{\Lambda}}y \iff & \big(\forall \lambda \in \tilde{\Lambda}, \forall z \in X : [yP_\lambda z \implies xP_\lambda z] \text{ and } [zP_\lambda x \implies zP_\lambda y]\big) \\
& \text{and} \\
& \big(\exists \lambda_0 \in \tilde{\Lambda}, \exists z \in X : [xP_{\lambda_0} z \text{ and } y \not P_{\lambda_0} z] \text{ or } [zP_{\lambda_0} y \text{ and } z \not P_{\lambda_0} x]\big).
\end{aligned}
$$

Then, for every $\lambda \in \tilde{\Lambda}$,

$$P_\lambda \subseteq T_{\tilde{\Lambda}} \subseteq P_{\overline{\lambda}}.$$

In particular, if $\overline{\lambda} \in \tilde{\Lambda}$ then $T_{\tilde{\Lambda}} = P_{\overline{\lambda}}$.

Proof. Fix a nonempty $\tilde{\Lambda} \subseteq \Lambda$.

We will first show that $P_\lambda \subseteq T_{\tilde{\Lambda}}$ for each $\lambda \in \tilde{\Lambda}$. Let $\lambda \in \tilde{\Lambda}$ and suppose $xP_\lambda y$. Because $P_\lambda \subseteq P_{\overline{\lambda}}$, we also have $xP_{\overline{\lambda}}y$, hence $x \geq y$.

We verify the two clauses in the definition of $T_{\tilde{\Lambda}}$.

Fix any $\lambda' \in \tilde{\Lambda}$ and any $z \in X$. If $yP_{\lambda'}z$, apply ordered discovery OD with (x,z,y) replaced by (x,y,z): since $x \succeq y$ and $yP_{\lambda'}z$, it follows that $xP_{\lambda'}z$. If $zP_{\lambda'}x$, apply ordered discovery OD directly: since $x \succeq y$ and $zP_{\lambda'}x$, it follows that $zP_{\lambda'}y$.

Take $\lambda_0 := \lambda$ and $z := y$. Then $xP_{\lambda_0}z$ holds. Moreover, by irreflexivity, $y\not{P}_{\lambda_0}y$. Thus the existential witness clause holds, and $xT_{\tilde{\Lambda}}y$.

Next, we show that $T_{\tilde{\Lambda}} \subseteq P_{\overline{\lambda}}$. Let $xT_{\tilde{\Lambda}}y$. We show $xP_{\overline{\lambda}}y$.

Suppose not. Then $x\not{P}_{\overline{\lambda}}y$, which by definition of the benchmark weak order implies $y \succeq x$.

Let (λ_0, z) be a witness for the existential clause in the definition of $T_{\tilde{\Lambda}}$.

If $xP_{\lambda_0}z$ and $y\not{P}_{\lambda_0}z$, then ordered discovery OD with $y \succeq x$ implies $yP_{\lambda_0}z$, contradicting $y\not{P}_{\lambda_0}z$.

If $zP_{\lambda_0}y$ and $z\not{P}_{\lambda_0}x$, then ordered discovery OD with $y \succeq x$ implies $zP_{\lambda_0}x$, contradicting $z\not{P}_{\lambda_0}x$.

Thus $y \succeq x$ is impossible, and so $xP_{\overline{\lambda}}y$, as desired.

Finally, if $\overline{\lambda} \in \tilde{\Lambda}$, then the left inclusion gives $P_{\overline{\lambda}} \subseteq T_{\tilde{\Lambda}}$ while the right inclusion gives $T_{\tilde{\Lambda}} \subseteq P_{\overline{\lambda}}$, hence equality. ■

Corollary 1. *Let $\Lambda' \subseteq \Lambda'' \subseteq \Lambda$ be nonempty sets of observed cognition indices. Then*

$$T_{\Lambda'} \subseteq T_{\Lambda''}.$$

In particular, enlarging the set of observed cognition indices can only expand the revealed relation $T_{\tilde{\Lambda}}$.

Proof. Fix $x, y \in X$ with $xT_{\Lambda'}y$. By Proposition 4, $xP_{\overline{\lambda}}y$ and hence $x \succeq y$.

We verify that $xT_{\Lambda''}y$.

Let $\lambda \in \Lambda''$ and $z \in X$. If $yP_\lambda z$, then $xP_\lambda z$ follows from ordered discovery OD using $x \succeq y$. If $zP_\lambda x$, then $zP_\lambda y$ follows from ordered discovery OD using $x \succeq y$. Thus the universal clause holds for every $\lambda \in \Lambda''$.

Because $xT_{\Lambda'}y$, there exists $\lambda_0 \in \Lambda'$ and $z_0 \in X$ satisfying the witness clause. Since $\Lambda' \subseteq \Lambda''$, the same (λ_0, z_0) witnesses the existential clause for $T_{\Lambda''}$.

Therefore $xT_{\Lambda''}y$ and so $T_{\Lambda'} \subseteq T_{\Lambda''}$. ■

Lemma 1. *For each $\lambda \in \Lambda$ and all $x, y \in X$,*

$$x\,\mathrm{TC}\big(\succeq_\lambda^{\mathrm{wk}}\big)\,y \quad\iff\quad \Big(\forall z \in X : yP_\lambda z \implies xP_\lambda z\Big) \text{ and } \Big(\forall z \in X : zP_\lambda x \implies zP_\lambda y\Big).$$

Proof. For $\succeq_\lambda^{\mathrm{wk}}$, we have $\uparrow_{\succeq_\lambda^{\mathrm{wk}}}(y) = \{z : z\not{P}_\lambda y\}$ and $\downarrow_{\succeq_\lambda^{\mathrm{wk}}}(x) = \{z : x\not{P}_\lambda z\}$. The condition $\uparrow(y) \subseteq \uparrow(x)$ is equivalent (by contraposition) to $\forall z : zP_\lambda x \implies zP_\lambda y$. Likewise, $\downarrow(x) \subseteq \downarrow(y)$ is equivalent to $\forall z : yP_\lambda z \implies xP_\lambda z$. ■

Proposition 5. *For every nonempty $\tilde{\Lambda} \subseteq \Lambda$,*

$$T_{\tilde{\Lambda}} = \succ_{\tilde{\Lambda}} = \{(x,y) : x \succsim_{\tilde{\Lambda}} y \text{ and } y \not\succsim_{\tilde{\Lambda}} x\}.$$

In particular, $T_{\tilde{\Lambda}}$ is transitive and irreflexive (a strict partial order).

Proof. By Lemma 1, the universal dominance clause in the definition of $T_{\tilde{\Lambda}}$ is exactly $x\,\mathrm{TC}(\succeq_{\lambda}^{\mathrm{wk}})\,y$ for every $\lambda \in \tilde{\Lambda}$, i.e. $x \succsim_{\tilde{\Lambda}} y$.

The witness clause in $T_{\tilde{\Lambda}}$ says that for some $\lambda_0 \in \tilde{\Lambda}$, the reverse dominance fails at level λ_0. By Lemma 1 again, this is equivalent to $y \not\mathrel{\mathrm{TC}(\succeq_{\lambda_0}^{\mathrm{wk}})} x$, which in turn is equivalent to $y \not\succsim_{\tilde{\Lambda}} x$.

Therefore $xT_{\tilde{\Lambda}}y$ iff $x \succsim_{\tilde{\Lambda}} y$ and not $y \succsim_{\tilde{\Lambda}} x$, i.e. iff (x,y) lies in the asymmetric part of $\succsim_{\tilde{\Lambda}}$.

Since $\succsim_{\tilde{\Lambda}}$ is an intersection of preorders, it is a preorder, and its asymmetric part is a strict partial order. Hence $T_{\tilde{\Lambda}}$ is transitive and irreflexive. ∎