Universitat Pompeu Fabra Barcelona
Department of Economics and Business

Economics Working Paper Series

Working Paper No. 1913

Random preference model

Mohammad Ghaderi, Kamel Jedidi,
Miłosz Kadziński, and Bas Donkers

July 2025

Random Preference Model

Mohammad Ghaderi,[a,b,c] Kamel Jedidi,[d], Miłosz Kadziński,[e], Bas Donkers[f]

[a]Department of Economics and Business, Pompeu Fabra University, [b]Barcelona School of Economics, [c]Barcelona School of Management, mohammad.ghaderi@upf.edu; [d]Columbia Business School, Columbia University, kj7@gsb.columbia.edu; [e]Institute of Computing Science, Poznan University of Technology, milosz.kadzinski@cs.put.poznan.pl; [f]Erasmus School of Economics, Erasmus University Rotterdam, donkers@ese.eur.nl

Abstract. We introduce the Random Preference Model (RPM), a non-parametric and flexible discrete choice model. RPM is a rank-based stochastic choice model where choice options have multi-attribute representations. It takes preference orderings as the main primitive and models choices directly based on a distribution over partial or complete preference orderings over a finite set of alternatives. This enables it to capture context-dependent behaviors while maintaining adherence to the regularity axiom. In its output, it provides a full distribution over the entire preference parameter space, accounting for inferential uncertainty due to limited data. Each ranking is associated with a subspace of utility functions and assigned a probability mass based on the expected log-likelihood of those functions in explaining the observed choices. We propose a two-stage estimation method that separates the estimation of ranking-level probabilities from the inference of preference parameters variation for a given ranking, employing Monte Carlo integration with subspace-based sampling. To address the factorial complexity of the ranking space, we introduce scalable approximation strategies: restricting the support of RPM to a randomly sampled or orthogonal basis subset of rankings and using partial permutations (top-k lists). We demonstrate that RPM can effectively recover underlying preferences, even in the presence of data inconsistencies. The experimental evaluation based on real data confirms RPM variants consistently outperform multinomial logit (MNL) in both in-sample fit and holdout predictions across different training sizes, with support-restricted and basis-based variants achieving the best results under data scarcity. Overall, our findings demonstrate RPM's flexibility, robustness, and practical relevance for both predictive and explanatory modeling.

Funding: M. Ghaderi's work was supported by Ramón y Cajal grant, funded by the Spanish Ministry of Science and Innovation [RYC2021-033938-I], and the Severo Ochoa Programme for Centres of Excellence in R&D [CEX2024-001476-S].

Key words: choice models, nonparametric modeling, rankings, context-dependent preference, random utility

1. Introduction

Modeling preferences from observed choices is a pivotal theme in many disciplines and plays a central role in mainstream economics (Samuelson 1938, Richter 1966, Afriat 1967, McFadden 2001, Chambers and Echenique 2016), decision theory (Fishburn 1970), operations research (Ben-Akiva and Lerman 1985, Train 2009, Farias et al. 2013), marketing (Louviere et al. 2000, Toubia et al. 2003), and psychology (Lichtenstein and Slovic 2006, Bettman et al. 1998), with wide-ranging applications in welfare analysis, policy design, healthcare, product development, pricing, and transportation (de Bekker-Grob et al. 2012, Hensher 1994, Ben-Akiva and Bierlaire 1999, Miller et al. 2011). Traditional approaches model preferences by assigning real-valued scores – commonly called *utilities* – to alternatives. These models begin with a utility function, shaped by

distributional and functional form assumptions, and map utility vectors to choice probabilities (Luce 1959, Feng et al. 2017). Representative models include attraction models (Bell et al. 1975, Gallego et al. 2015), random utility models (RUMs) (McFadden 1973, Train 2009), exponomial choice models (Alptekinoğlu and Semple 2016, Aouad et al. 2023), and process-based models such as elimination by aspects (Tversky 1972), preference trees (Tversky and Sattath 1979), and Markov chain choice models (Blanchet et al. 2016, Feldman and Topaloglu 2017, Désir et al. 2024).

While these utility-based frameworks have been highly influential, they face critical limitations. Epistemically, they define preferences as summaries of observed behavior, thus embedding context effects implicitly rather than modeling them directly. This can obscure systematic deviations from rational choice axioms, such as trade-off contrast effects (Simonson and Tversky 1992), extremeness aversion (Simonson and Tversky 1992), or salience-driven behavior (Bordalo et al. 2013). Technically, the parametric assumptions underpinning utility models risk misspecification and reduce flexibility in high-dimensional or heterogeneous choice environments (Abdellaoui 2000, Bleichrodt and Pinto 2000, Ghaderi and Kadziński 2021).

Rank-based choice models provide a compelling alternative. These nonparametric approaches characterize decision-makers (DMs) through partial or complete preference orderings over a finite set of alternatives (Farias et al. 2013, van Ryzin and Vulcano 2015, 2017, Bertsimas and Mišić 2019, Sturt 2025). In contrast to utility-based models, they adopt preference orderings as the fundamental primitive, directly generating choice behavior. This methodological shift enables rank-based models to flexibly capture behaviors that violate standard axioms such as independence from irrelevant alternatives (IIA) (Luce 1959) or order-independence (Simonson and Tversky 1992), and to represent empirically observed phenomena like the attraction effect, asymmetric dominance, and compromise effects (Huber et al. 1982, Tversky and Simonson 1993, Kivetz et al. 2004) without relying on parametric assumptions.

Although the theoretical foundations of rank-based models date back decades (Block and Marschak 1959, Falmagne 1978, McFadden and Richter 1990), their appeal has grown in recent years due to their intuitive structure, modeling flexibility, and empirical realism (Farias et al. 2013, Jagabathula 2014, Bertsimas and Mišic 2015). These models subsume broad classes of RUMs but avoid the rigidity and specification risks of traditional methods. Nonetheless, their adoption has been hindered by the combinatorial complexity of the preference space and the computational cost of estimation (van Ryzin and Vulcano 2017, Aouad et al. 2021).

In this paper, we develop a general operational framework for rank-based choice models that extends their applicability to contexts where alternatives are described by multiple attributes. This is particularly relevant for analyzing consumer substitution patterns in domains such as product design, pricing, and conjoint analysis (Toubia et al. 2007, Louviere et al. 2011). Our model supports a wide range of outputs, including detailed measures of inferential uncertainty, making it suitable not only for predictive tasks, such as revenue forecasting and assortment optimization, but also for explanatory and normative purposes in marketing, economics, and social choice (Baldiga and Green 2013).

The proposed approach builds upon the random preference literature (Block and Marschak 1959, Barberá and Pattanaik 1986, Falmagne 1978, Fishburn and Falmagne 1989), incorporating recent developments in distributionally robust, nonparametric inference (Srinivasan and Shocker 1973, Farias et al. 2013, Ghaderi et al. 2017, Ghaderi 2017, Farias et al. 2020, Ghaderi and Kadziński 2021). It captures two distinct sources of uncertainty: inferential uncertainty due to limited or noisy data, and behavioral stochasticity arising from cognitive processes such as inattention, inertia, or contingent reasoning (Samuelson and Zeckhauser 1988, Salant and Rubinstein 2008, Kadziński et al. 2020). While these uncertainties differ in interpretation, they are often empirically indistinguishable, and a probabilistic framework provides a unified treatment for both.

To address the computational challenges associated with rank-based modeling, we introduce a suite of approximation strategies that enable users to trade off model fidelity and estimation complexity. We empirically evaluate our model using both synthetic and real-world data, demonstrating its practical accuracy, robustness, and interpretability.

2. Primitives

Let X be a finite set, where each element $x \in X$ represents a choice option, possibly including a no-choice option. Denote by $\mathcal{X}$ the collection of all nonempty subsets of X, where each $S \in \mathcal{X}$ represents a choice set, i.e., a menu or choice task.[1] In each choice task, the DM, when faced with a menu $S \in \mathcal{X}$, selects an alternative $c(S)$, where c is a *choice function*, i.e., a function from $\mathcal{X}$ to X satisfying $c(S) \in S$.

The standard model of choice proceeds by positing a binary relation that represents the DM's preferences, reflects an inherent decision-making capacity, or serves as a mathematical construct approximating observed behavior under an *as if* argument. This binary relation is typically assumed

[1] We consider menus of size at least two, since choices from singletons are trivial.

to be complete and transitive – properties that together define a *rational preference relation*. A rational preference relation induces a ranking over X, which in turn determines the choice function c for any menu $S \in \mathcal{X}$.

Let $\mathcal{P}$ denote the set of all complete, transitive, and asymmetric binary relations on X. Then, $|\mathcal{P}| = |X|!$. A substantial body of economic theory focuses on identifying conditions – such as the Weak Axiom of Revealed Preference – that allow the construction of a rational preference relation $P \in \mathcal{P}$ that determines outcomes of the choice function c. When such a relation exists, c is said to be *rationalized*. However, the axioms that lead to a rational preference relation are rarely satisfied in real-world data, where individuals often make inconsistent choices from identical menus (Agranov and Ortoleva 2017), or where choice data are observed at the population level in the form of choice shares of options within a menu across individuals. Modeling such data requires a joint probability distribution over $X \times \mathcal{X}$, representing a *stochastic choice function* (SCF), which extends the deterministic function c.

The randomness of such SCFs admits multiple interpretations. At the population level, it can represent taste heterogeneity among individuals. At the individual level, choices may appear *observationally stochastic* due to limited information or inferential uncertainty from the analyst's perspective, measurement errors, or *intrinsically stochastic* behavior arising from cognitive factors such as inattention, inertia, or intentional randomization—as modeled in perturbed utility frameworks. While interpretations vary across domains, the formalism and operationalization of SCFs remain consistent.

Definition 1 (Stochastic Choice Function). A stochastic choice function is a mapping $\rho : X \times \mathcal{X} \to [0, 1]$ such that $\rho(x, S) = 0$ for all $x \notin S$, and $\sum_{x \in S} \rho(x, S) = 1$.

SCFs are natural modeling tools for observed choice behavior and are amenable to empirical validation. A *stochastic choice model* (SCM) is a parameterized representation of an SCF, where different parameterizations lead to distinct choice models. For example, suppose there exists a real-valued function $v : X \to \mathbb{R}$ such that

$$\rho(x, S) = \frac{v(x)}{\sum_{y \in S} v(y)} \quad \text{for all } (x, S) \in X \times \mathcal{X}.$$

This yields the well-known *Luce choice model* (Luce 1959). If $v(x) \in \mathbb{R}_{++}$ and we apply a logarithmic transformation $u(x) = \ln v(x)$, the model becomes the *logit model* (McFadden 1973, 2001). Note that the positivity assumption ($v(x) > 0$) implies that $\rho(x, S) > 0$, an assumption not empirically

testable, regardless of data volume. An SCF is said to have a Luce representation if such a positive function v exists.

For binary menus, if there exists a real function u and a strictly increasing function ϕ such that $\rho(x, \{x, y\}) = \phi(u(x) - u(y))$, the SCF is said to follow a *Fechnerian model*. Specific choices of ϕ, such as the logistic or Gaussian CDF, yield the binary logit and probit models, respectively.

A more general class of SCMs is captured by the *random utility model* (RUM). It posits a random vector $\tilde{U}$ over alternatives such that

$$\rho(x, S) = \mathbb{P}\left\{\tilde{U}_x = \max_{y \in S} \tilde{U}_y\right\},$$

where $\mathbb{P}$ is a probability measure. Defining $u(x) = \mathbb{E}[\tilde{U}_x]$ and $\varepsilon_x = \tilde{U}_x - u(x)$, and assuming *positivity*, the model can be expressed as:

$$\rho(x, S) = \mathbb{P}\left\{\varepsilon_x \geq \max_{y \in S}(u(y) - u(x) + \varepsilon_y)\right\},$$

a formulation known as the *discrete choice model* (Ben-Akiva and Lerman 1985). Varying the joint distribution of the ε_x's yields various SCMs (Train 2009). For example, assuming independent ε_x following an Extreme Value Type-I distribution recovers the multinomial logit model (MNL) (McFadden 1973, 2001). Relaxing the independence assumption and introducing a block covariance structure yields the nested logit model (Ben-Akiva 1973). For binary menus with identically distributed ε_x, the RUM reduces to a Fechnerian model.

Each of the above SCMs imposes specific conditions on the SCF they aim to represent. For instance, the Luce model assumes both *positivity* and the Luce *independence* property. In fact, an SCF has a Luce representation if and only if it satisfies these two conditions. Characterizing choice models through properties of the SCF has the advantage of yielding assumptions that are, in principle, directly testable using observed choice data. However, verifying such assumptions often requires extensive data, particularly observations across the full range of menus in $\mathcal{X}$. As a result, they may not be refutable with limited or sparse data.

Moreover, in real-world applications where there is a sufficient variation in the menus presented, these assumptions are rarely satisfied. For instance, there is overwhelming evidence of violations of the independence assumption (Huber et al. 1982, Tversky and Simonson 1993). Consequently, the analyst's task becomes selecting a choice model that best *approximates* the DM's preferences. Still, understanding the theoretical properties of SCMs in an idealized setting – where the full SCF is known – can guide the selection of an appropriate model. For this reason, we briefly review some of

the most commonly used assumptions in mainstream SCMs before presenting the technical details of our proposed model.

Axiom 1 (Luce Independence) *For any $x, y \in S \subset T \in \mathcal{X}$, it holds that $\frac{\rho(x,S)}{\rho(y,S)} = \frac{\rho(x,T)}{\rho(y,T)}$.*

The Luce independence assumption, also known as the IIA assumption, asserts that the odds of choosing x over y are invariant to the presence of other alternatives. A weaker version of IIA is order independence, requiring that the ordinal ranking of two options remains unchanged across menus:

Axiom 2 (Order Independence) *For any $x, y \in S \subset T \in \mathcal{X}$,*

$$\rho(x,S) - \rho(y,S) \geq 0 \iff \rho(x,T) - \rho(y,T) \geq 0.$$

For binary menus, both the Luce and logit models are special cases of Fechnerian models. When $|X| \leq 3$, Block and Marschak (1959) showed that the following condition is necessary and sufficient for RUM representability:

Axiom 3 (Regularity) *For any $x \in S \subset T \in \mathcal{X}$, $\rho(x,S) \geq \rho(x,T)$.*

Regularity, also called monotonicity, ensures that adding alternatives to a menu does not increase the choice probability of an existing option. It is a central axiom in economic choice theory and is often a necessary condition even in models allowing boundedly rational behavior. However, for $|X| > 3$, regularity is necessary but not sufficient for RUM representability. For example, when $|X| = 4$, Block and Marschak (1959) identified the need for an additional axiom:

Axiom 4 (Supermodularity) *For any three disjoint sets $S, T, V \in \mathcal{X}$ and $x \in S$,*

$$\rho(x, T \cup S) - \rho(x, T \cup S \cup V) \leq \rho(x,S) - \rho(x, S \cup V).$$

This axiom implies diminishing marginal losses in choice probability as menus expand (Strzalecki 2017). Together, regularity and supermodularity define decreasing and convex properties of $\rho(x,S)$ across partially ordered menus. In other words, supermodularity and regularity define properties of $\rho(x,S)$ when viewed as a function of all partially ordered menus containing x. Regularity assumes that this function is decreasing. Supermodularity assumes that it is convex. When menus are binary, an alternative condition is sufficient:

Axiom 5 (Triangle Inequality) *For any $x, y, z \in X$,*

$$\rho(x, \{x, z\}) \leq \rho(x, \{x, y\}) + \rho(y, \{y, z\}).$$

The triangle inequality is both necessary and sufficient for RUM representability when menus are binary and $|X| < 6$ (Block and Marschak 1959, Fishburn 1992), but insufficient when $|X| \geq 6$ (McFadden and Richter 1970). In general, for arbitrary menu sizes, RUM representability requires nonnegative Block-Marschak polynomials (Barberá and Pattanaik 1986, Falmagne 1978). However, verifying this condition empirically is infeasible with limited data.

3. Random Preference Model

This section introduces the *random preference model* (RPM). While it satisfies *regularity* – a foundational axiom in economic theory – it relaxes stronger assumptions such as positivity and Luce independence. We demonstrate that RPM effectively captures context-dependent behaviors, including the asymmetric dominance and compromise effect, and accommodates violations of the IIA and order independence. This challenges the widespread belief that effects like asymmetric dominance necessarily imply violations of regularity. We argue that regularity is violated only in extreme cases, which are rarely encountered in practice (Frederick et al. 2014). The regularity-consistent structure of RPM, combined with its flexibility in modeling moderate context effects, makes it a compelling and practically useful choice model.

The RPM defines a probability distribution over $\mathcal{P}$, the set of all complete, transitive, and asymmetric binary relations on X. Let $\Delta(\mathcal{P})$ denote the $|X|$-dimensional simplex of all probability distributions over $\mathcal{P}$. From the analyst's perspective, each probability in this vector represents the likelihood that a given preference relation governs the observed choice behavior.

Definition 2 (Random Preference Model). A random preference model is a pair $(\mu, \mathcal{P})$, where $\mathcal{P}$ is the set of complete, transitive, and asymmetric binary relations over the set of choice options X, and $\mu \in \Delta(\mathcal{P})$ is a probability distribution over the elements of $\mathcal{P}$.

We now consider how the RPM describes choice behaviors. For an observed choice (x, S), define:

$$Q(x, S) = \{P \in \mathcal{P} : xPy \text{ for all } y \in S, y \neq x\}, \tag{1}$$

where xPy means $(x, y) \in P$, i.e., x is preferred to y under preference P, or equivalently, x precedes y in the ranking induced by P.

Lemma 1. *For any $S \in \mathcal{X}$, $\mu\big(Q(x, S)\big) \geq 0$ and $\sum_{x \in S} \mu\big(Q(x, S)\big) = 1$.*

Proof sketch. For any $S \in \mathcal{X}$, the sets $Q(x,S)$ form a partition of $\mathcal{P}$. First, for $x \neq z \in S$, we have $Q(x,S) \cap Q(z,S) = \emptyset$, because if $P \in Q(x,S) \cap Q(z,S)$, then xPz and zPx, contradicting the asymmetry of P. Second, since every $P \in \mathcal{P}$ ranks all elements in S, it has a unique maximal element in S due to completeness, transitivity, and asymmetry (Rubinstein 2012). Thus, $\bigcup_{x \in S} Q(x,S) = \mathcal{P}$, and since μ is a probability measure over $\mathcal{P}$, the result follows.

We say that an SCF ρ is *induced* by an RPM μ – denoted ρ_μ – if there exists $\mu \in \Delta(\mathcal{P})$ such that:

$$\rho(x,S) = \mu\big(Q(x,S)\big). \tag{2}$$

That is, the probability of choosing x from S equals the total probability assigned to all preference relations in $\mathcal{P}$ that rank x above all other elements in S. Equivalently, $\rho_\mu(x,S) = \sum_{P \in \mathcal{P}} \mu(P) \times \mathbb{I}[xPy \text{ for all } y \in S \setminus \{x\}]$, where $\mathbb{I}$ is the indicator function.

Every RPM induces a unique SCF, but the converse does not generally hold. An SCF may admit multiple distinct RPM representations, even with disjoint supports (Fishburn 1998). This identification issue is common in stochastic choice modeling. Nonetheless, when the goal is to identify a modal preference or representative agent – often desirable due to convexity considerations – RPM offers advantages over utility-based representations. Full identification of an RPM requires domain restrictions (Turansick 2022, Apesteguia et al. 2017). While preference relations may be recoverable under suitable assumptions (Chambers et al. 2018), identifying utility functions is generally much more difficult. Since utility is ordinal, its distribution cannot be identified from finite choice data; at best, we can hope to recover a distribution over preferences (Strzalecki 2017). The following result alleviates identification concerns for prediction tasks:

THEOREM 1 **(Falmagne (1978))**. *If μ and μ' are two RPMs that represent the same SCF ρ, then the probability of any alternative $x \in X$ being ranked in position k is the same under both μ and μ'.*

In real-world applications with noisy or limited data, no single RPM or RUM can perfectly explain observed choices. The analyst's objective, therefore, becomes finding a model that adequately *approximates* the underlying preferences. In this regard, RPM has notable advantages over utility-based models. Its nonparametric structure and inherent flexibility make it well-suited for capturing a wide range of boundedly rational behaviors, including violations of IIA, order independence, and context-dependent effects such as the compromise effect and trade-off contrast. We illustrate these phenomena using a simple running example before presenting the multiattribute extension of RPM. For binary menus, we abbreviate $\rho(x,\{x,y\})$ as $\rho(x,y)$.

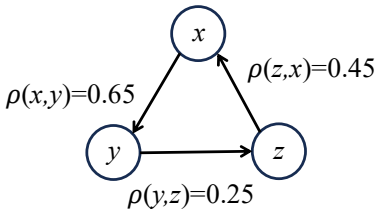

Figure 1 **The system of binary choice probabilities in Example 1.**

EXAMPLE 1. Consider the choice set $X = \{x, y, z\}$ and the following system of binary choice probabilities: $\rho(x, y) = 0.65$, $\rho(y, z) = 0.25$, and $\rho(z, x) = 0.45$. The probability $\rho(x, \{x, y, z\})$ is unknown. The structure of the binary choice probabilities is illustrated in Figure 1.

Table 1 enumerates all six possible complete rankings over the elements of X. Each vector $(\mu_1, \ldots, \mu_6)$ defines an RPM, where $\mu_n = \mu(P_n)$ represents the probability mass assigned to ranking P_n. For instance, consider the distribution $\mu = (0, 0.40, 0.15, 0.10, 0.25, 0.10)$, which is fully consistent with the observed system of binary choice probabilities. To verify this, observe that:

$$\rho_\mu(x, y) = \mu_1 + \mu_2 + \mu_5 = 0 + 0.40 + 0.25 = 0.65,$$

which matches the given value of $\rho(x, y)$. Similar calculations confirm the consistency for the other two binary menus. For the ternary menu $\{x, y, z\}$, we have $\rho_\mu(x, \{x, y, z\}) = 0.40$. Note that this does not violate the regularity axiom, as $\rho_\mu(x, \{x, y, z\})$ is less than both $\rho(x, y) = 0.65$ and $\rho(x, z) = 1 - \rho(z, x) = 0.55$. In fact, any SCF induced by an RPM satisfies regularity, as stated by the following lemma.

LEMMA 2. *Any SCF ρ induced by an RPM μ satisfies regularity. That is,*

$$\rho_\mu(x, S) \leq \min_{\substack{T \subset S \\ x \in T}} \rho_\mu(x, T).$$

The following discussion highlights three key behavioral phenomena – Luce independence, the compromise effect, and order dependence – and demonstrates how the RPM relates to each.

Table 1 **The population of rankings in Example 1.**

$\mathcal{P}$	P_1	P_2	P_3	P_4	P_5	P_6
μ	μ_1	μ_2	μ_3	μ_4	μ_5	μ_6
	x	x	y	y	z	z
Rankings:	y	z	x	z	x	y
	z	y	z	x	y	x

• **Luce Independence:** In the example, we observe that the ratio of choice probabilities between alternatives x and y varies across different menus, i.e.:

$$\frac{\rho_\mu(x,y)}{\rho_\mu(y,x)} = \frac{0.65}{0.35} > \frac{\rho_\mu(x,\{x,y,z\})}{\rho_\mu(y,\{x,y,z\})} = \frac{0.40}{0.25},$$

demonstrating that RPM does not satisfy IIA. In utility-based choice models, the nested logit model is one of the most widely used approaches for capturing violations of IIA (Ben-Akiva 1973, McFadden 1977). However, implementing nested logit requires the analyst to specify a nesting structure – a task that is subjective, often cumbersome, and computationally complex due to its combinatorial nature (Aboutaleb et al. 2020). Kohli and Jedidi (2017) showed that nested logit is a special case of the probabilistic model known as Elimination by Aspects (EBA) (Tversky 1972). Notably, RPM subsumes EBA (Bertsimas and Mišić 2019), and therefore also subsumes nested logit. Consequently, RPM is capable of modeling non-IIA choice data without requiring any nesting assumptions.

• **Compromise Effect:** The compromise effect arises from a cognitive phenomenon known as *trade-off contrast*, in which DM's choices are influenced by the presence of other options that imply different exchange rates between attributes such as quality and price. This contrast does not require any option to dominate another. However, when a menu includes an asymmetrically dominated decoy, the market share of the dominating option *relative* to its competitor tends to increase (Huber et al. 1982, Simonson and Tversky 1992). In such cases, the mere presence of a dominated decoy makes one option appear more attractive than another.

To quantify the compromise effect, one may adopt the metric proposed by Simonson and Tversky (1992) and Kivetz et al. (2004). Specifically, the effect of adding an option z to the menu $\{x,y\}$ on the choice probability of y is measured by:

$$D_z(y;x) = \frac{\rho(y,\{x,y,z\})}{\rho(y,\{x,y,z\}) + \rho(x,\{x,y,z\})} - \rho(y,x),$$

capturing the change in the *relative* popularity of y compared to x due to the presence of z. If both $D_z(y;x)$ and $D_x(y;z)$ are positive, one can conclude that y is a compromise option in the $\{x,y,z\}$ menu. In the case of asymmetric dominance, if y is clearly superior to z but x is not, hence z being a decoy to y, then Simonson and Tversky (1992) predicts that $D_z(y;x) > 0$, violating independence and potentially order independence axioms. However, it does not necessarily violate regularity, as we discuss later in this section.

The compromise effect is both a widely used strategy in menu design and a highly debated form of context-dependent preference (Kivetz et al. 2004). Frederick et al. (2014), through a series of 38 studies, failed to replicate the effect under various conditions and suggested that it may be confined to laboratory settings and stylized products with numeric attributes. In contrast, Wu and Cosguner (2020) found the effect to be empirically significant, conditional on the DM detecting the decoy – though detection rates were low. Simonson (2014) responded to these critiques by noting that the effect is often not observed when stronger choice drivers are at play, but it does tend to manifest, especially when price is among the attributes. Taken together, the literature suggests that while the compromise effect is robust, its magnitude is often more modest than originally claimed in popular accounts such as (Ariely 2008), which reported an increase in market share from 32% to 84% due to a decoy.

In our example, the change in the *relative* popularity of y compared to x or z due to the presence of, respectively, z or x can be quantified as: $D_z(y) = \frac{0.25}{0.25+0.40} - 0.35 = 0.035 > 0$ and $D_x(y) = \frac{0.25}{0.25+0.35} - 0.25 = 0.167 > 0$, indicating a discernible compromise effect for the RPM, with y being perceived as the middle option between x and z.

A common misconception in economics, marketing, and operations (Frederick et al. 2014, Strzalecki 2017, Berbeglia and Venkataraman 2023) is that the compromise effect necessarily violates the regularity axiom – the principle that adding alternatives should not increase the probability of choosing an existing option. However, as our example shows, this is not necessarily the case. The decoy's role is to enhance the relative appeal of one option without necessarily increasing its absolute share. Violations of regularity occur only under extreme conditions, such as when the decoy is never chosen (a case deemed uninteresting by Huber et al. 1982) or when shifts in relative share are excessively large (Simonson and Tversky 1992). Empirical studies suggest such violations are rare in practice (Frederick et al. 2014). Thus, RPM is capable of capturing realistic compromise effects while preserving regularity.

• **Order-Dependent Preferences:** The RPM shares the flexibility of models such as preference trees designed to accommodate violations of order independence (Kohli and Jedidi 2017). To illustrate, consider the RPM defined by $\mu = (\mu_1, \ldots, \mu_6) = (0, 0.30, 0.25, 0, 0.35, 0.10)$. This distribution is fully consistent with the binary choice system in Example 1, i.e., $\rho_\mu \equiv \rho$ for binary menus. For the two menus $S = \{x, z\}$ and $T = \{x, y, z\}$, we observe: $\rho(x, S) - \rho(z, S) = 0.55 - 0.45 = 0.10 > 0$, but $\rho(x, T) - \rho(z, T) = 0.30 - 0.45 = -0.15 < 0$, which violates the order independence axiom. From this, we conclude:

LEMMA 3. *The RPM can capture choice data when order independence does not hold.*

3.1. RPM: The Sweet Spot Between Rank- and Utility-Based Choice Models

Similar to the approaches of Barberá and Pattanaik (1986) and McFadden and Richter (1990), our RPM framework directly models preferences via a probability distribution over the set of strict linear orderings (rankings) of alternatives, rather than positing an underlying utility function. When the choice set X is finite, the RPM representation is equivalent to the RUM in terms of the induced choice probabilities (Block and Marschak 1959, McFadden 2001).

THEOREM 2 **(Block and Marschak (1959))**. *An SCF ρ is RPM-representable if and only if it is RUM-representable.*

This equivalence holds in the absence of structural assumptions on either representation. Although RPM and RUM yield the same SCF, their underlying parameters – and hence their interpretations – differ fundamentally. The RPM representation offers a simplified and transparent framework, as only ordinal information (i.e., the ranking of alternatives) is relevant to observed choices (Barberá and Pattanaik 1986). Since the two approaches are inherently different in the way they arrive at an SCF, differences in representations are consequential for empirical modeling and policy analysis, especially in settings where model parameters themselves are objects of interest, such as decision analysis and discrete choice applications involving alternatives evaluated in terms of multiple attribute (Fishburn 1970, McFadden 1974, Keeney and Raiffa 1993, Kadziński et al. 2012, Ghaderi and Kadziński 2021, Allen and Rehbeck 2023).

While RUM describes choices using a distribution over the utility space, and rank-based models describe choices based on a distribution over rankings, RPM describes choices based on a probability distribution over rankings where the probability of a given ranking is related to a distribution over the subspace of utility functions consisting of all utility functions compatible with this ranking. Hence, RPM relates RUM with rank-based choice models. On the one hand, preferences are modeled through rankings. On the other hand, given a ranking, RPM constructs an empirical distribution over the entire space of utility functions that represent it. This construction relates the ranking distribution to utility representations via a notion of expected log-likelihood, which is proportional to the ranking probability. Thus, for each ranking, we define a distribution over the compatible utility subspace such that the expected log-likelihood induced by that distribution aligns with the observed ranking probability. This mechanism enables utility-based inference while preserving the rank-based foundation of the model. Furthermore, discarding distributions over the utility subspaces, our framework also aligns with rank-based choice models, which are prominent

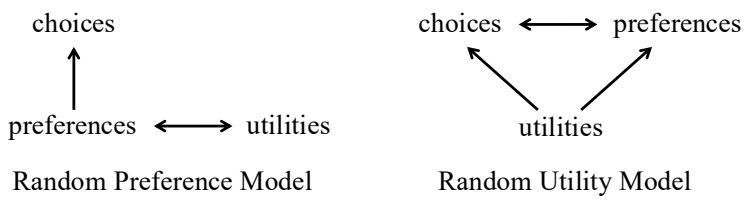

Figure 2 Relationship between preferences, choices, and utilities in RPM versus conventional RUM.

in the assortment optimization and revenue management literature, where alternatives are treated as nominal objects without a multi-attribute description.

Moreover, RPM and RUM differ in their epistemological underpinnings (Hausman 2011). RPM does not treat preferences as summaries of satisfaction or utility; rather, it considers preferences and rankings as primitive elements of behavior, which are reflected in satisfaction levels and represented by utility functions. In contrast, RUM typically takes utility as primitive, with preferences and choice behavior derived therefrom. These foundational differences also affect identifiability: RPM identifies preferences directly, while RUM identifies utilities as behavioral summaries.

In line with this perspective, we treat preferences as the primary construct generating choices, and utility functions as representational tools for these preferences. This stands in contrast to models that begin with utility as primitive and define choice probabilities directly from utility comparisons, thereby reducing preferences to mere summaries of choice behavior. Figure 2 illustrates this conceptual distinction.

In the next section, we outline the construction of RPM, how probabilities over rankings are constructed, and how the probability for each ranking is related to an empirically constructed distribution over the compatible utility subspace. Specifically, for each ranking, we construct a distribution over the corresponding utility subspace, such that the expected log-likelihood within each subspace matches the associated ranking probability in a structured and coherent manner. The structure arises from partitioning the utility space according to rankings – each subspace containing utility functions that represent a specific ranking. The coherence is ensured by enforcing that the expected log-likelihood within each subspace is consistent with the ranking's probability. This construction bridges the interpretive strengths of both RPM and RUM and facilitates utility-based inference grounded in ordinal choice behavior.

3.2. Construction of Random Preference Model

Let $\mathcal{U}$ denote the space of utility functions. A utility function $u \in \mathcal{U}$ is characterized by a parameter vector $\boldsymbol{\theta} \in \Theta$, and is said to represent a preference relation $P \in \mathcal{P}$ if $u(x;\boldsymbol{\theta}) > u(y;\boldsymbol{\theta})$ whenever

xPy. For a given preference relation $P \in \mathcal{P}$, define the subspace $\Theta^P \subseteq \Theta$ as the set of parameters corresponding to utility functions $\mathcal{U}^P \subseteq \mathcal{U}$ that represent P.

Remark 1. The collection of preference relations $\mathcal{P}$ partitions the parameter space Θ. That is, $\bigcup_{P\in\mathcal{P}} \Theta^P = \Theta$, and for any two distinct preference relations $P_m, P_n \in \mathcal{P}$, we have $\Theta^{P_m} \cap \Theta^{P_n} = \emptyset$.

Remark 2. For any preference relation $P \in \mathcal{P}$, the subspace Θ^P is convex regardless of the functional form of the utility functions. Specifically, if $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \Theta^P$, then any convex combination $t\boldsymbol{\theta} + (1-t)\boldsymbol{\theta}' \in \Theta^P$ for all $t \in [0,1]$.

We assert that the ability of a preference $P \in \mathcal{P}$ to explain the observed choice data is proportional to the collective ability of the utility functions $u \in \mathcal{U}^P$ in doing so. For a $\boldsymbol{\theta} \in \Theta^P$, we define its explanatory power with respect to the observed dataset $\mathcal{D} = \{(x, S)\}$ as follows:

$$\beta(\boldsymbol{\theta}) = \frac{1}{|\mathcal{D}|} \sum_{(x,S)\in\mathcal{D}} \log\Big(\frac{e^{u(x;\boldsymbol{\theta})}}{\sum_{y\in S} e^{u(y;\boldsymbol{\theta})}}\Big). \tag{3}$$

Here, $\beta(\boldsymbol{\theta})$ corresponds to the expected log-likelihood by which $\boldsymbol{\theta}$ chooses x from S for a randomly selected choice data (x, S). A higher value of $\beta(\boldsymbol{\theta})$ indicates stronger discriminatory power in distinguishing the chosen alternative from the unchosen ones in menu S, under a fixed preference relation P [2]. We link the probability mass assigned to each preference P to the expected value of this log-likelihood-based discrimination measure. Specifically, let f be an unknown probability distribution over the space of parameters Θ. Then, $\phi = \mathbb{E}_{f|P}\big[\beta(\boldsymbol{\theta})\big]$, where $f|P = f(\boldsymbol{\theta}|P)$ $f|P = f(\boldsymbol{\theta} \mid P)$ denotes a conditional probability given a specific ranking. Hence, we define

$$\phi = \int_{\boldsymbol{\theta}\in\Theta^P} \beta(\boldsymbol{\theta}) f(\boldsymbol{\theta}|P) d\boldsymbol{\theta} \tag{4}$$

and therefore

$$\mu(P) = \frac{e^{\sigma\phi}}{\sum_n e^{\sigma\phi_n}}, \tag{5}$$

where $\sigma \geq 0$ is a concentration parameter that governs the sharpness of the distribution μ. Intuitively, σ captures the explanatory strength of the variables involved in the utility specification. Therefore, it can serve as an objective measure to test their relevance. Larger values of σ indicate greater explanatory power, while $\sigma = 0$ yields a uniform distribution over preference orders P.

Using the same framework, we define the conditional distribution over the subspace Θ^P. Specifically, the probability density of a utility function $u(\cdot;\boldsymbol{\theta}) \in \mathcal{U}^P$ is defined as:

$$f(\boldsymbol{\theta}|P) = \frac{e^{\lambda\beta(\boldsymbol{\theta})}}{\int_{\Theta^P} e^{\lambda\beta(\boldsymbol{\theta})} d\boldsymbol{\theta}}, \tag{6}$$

[2] Utilities are always scaled in a $[0, 1]$ interval.

where $\lambda \geq 0$ is a concentration parameter for the distribution over each Θ^P that may vary across preference relations $P \in \mathcal{P}$. Substituting (6) into (4) yields:

$$\phi = \frac{\int_{\Theta^P} \beta(\boldsymbol{\theta}) e^{\lambda \beta(\boldsymbol{\theta})} d\boldsymbol{\theta}}{\int_{\Theta^P} e^{\lambda \beta(\boldsymbol{\theta})} d\boldsymbol{\theta}}. \tag{7}$$

Computing the value of ϕ involves integration over a potentially high-dimensional parameter space and a function $\beta(\boldsymbol{\theta})$ that lacks a closed-form expression. Nonetheless, ϕ can be approximated as a function of λ by sampling from Θ^P. In fact, it is possible to estimate ϕ without directly estimating λ by decomposing the estimation process into two stages – a procedure we describe in the next section.

4. Estimation of Random Preference Model

To estimate the parameters of the RPM from observed choice data, we consider a maximum likelihood and a distance-based estimation framework. Regardless of the estimation process, it is important to note that the parameter ϕ, defined in 4, corresponds to the expected log likelihood of the compatible utility subspace, even with the distance-based estimation.

Let $\{(x_t, S_t)\}_{t=1}^T$ denote a finite collection of observed choices. The corresponding likelihood function is defined as:

$$\mathcal{L}(\sigma, \{\phi_n\}_n, \{\lambda_n\}_n) = \prod_{t=1}^{T} \mu\Big(Q(x_t, S_t)\Big). \tag{8}$$

This likelihood function is generally nonlinear and nonconvex, making direct optimization computationally demanding. However, we demonstrate that the likelihood can be decomposed by isolating the λ parameters via bounding the ϕ terms. This facilitates a two-stage estimation procedure: we first solve for (σ, ϕ) under bounded constraints and subsequently recover λ by solving a system of linear equations.

LEMMA 4. *For any $P \in \mathcal{P}$, $\mathbb{E}_{uni(\Theta^P)}\big[\beta(\boldsymbol{\theta})\big] \leq \phi \leq \max_{\boldsymbol{\theta} \in \Theta^P} \beta(\boldsymbol{\theta})$.*

Proof The derivative of ϕ with respect to λ is given by:

$$\frac{d\phi}{d\lambda} = \frac{d}{d\lambda}\left(\frac{\int_{\Theta^P} \beta(\boldsymbol{\theta}) e^{\lambda \beta(\boldsymbol{\theta})}\, d\boldsymbol{\theta}}{\int_{\Theta^P} e^{\lambda \beta(\boldsymbol{\theta})}\, d\boldsymbol{\theta}}\right).$$

Applying the quotient rule, we obtain:

$$\frac{d\phi}{d\lambda} = \frac{\left(\int_{\Theta^P} \beta(\boldsymbol{\theta})^2 e^{\lambda \beta(\boldsymbol{\theta})}\, d\boldsymbol{\theta}\right)\left(\int_{\Theta^P} e^{\lambda \beta(\boldsymbol{\theta})}\, d\boldsymbol{\theta}\right) - \left(\int_{\Theta^P} \beta(\boldsymbol{\theta}) e^{\lambda \beta(\boldsymbol{\theta})}\, d\boldsymbol{\theta}\right)^2}{\left(\int_{\Theta^P} e^{\lambda \beta(\boldsymbol{\theta})}\, d\boldsymbol{\theta}\right)^2}.$$

But this is the variance of $\beta(\boldsymbol{\theta})$ under the distribution $f(\boldsymbol{\theta}|P)$:

$$\frac{d\phi}{d\lambda} = \frac{\int_{\Theta^P} \beta(\boldsymbol{\theta})^2 e^{\lambda\beta(\boldsymbol{\theta})}\, d\boldsymbol{\theta}}{\int_{\Theta^P} e^{\lambda\beta(\boldsymbol{\theta})}\, d\boldsymbol{\theta}} - \left(\frac{\int_{\Theta^P} \beta(\boldsymbol{\theta}) e^{\lambda\beta(\boldsymbol{\theta})}\, d\boldsymbol{\theta}}{\int_{\Theta^P} e^{\lambda\beta(\boldsymbol{\theta})}\, d\boldsymbol{\theta}}\right)^2$$

Under the distribution $f(\boldsymbol{\theta}|P)$ in Eq. (6), the first term is the expected value of squared $\beta(\boldsymbol{\theta})$ and the second term is the squared expected value of $\beta(\boldsymbol{\theta})$, that is:

$$\frac{d\phi}{d\lambda} = \mathbb{E}_f[\beta(\boldsymbol{\theta})^2] - \left(\mathbb{E}_f[\beta(\boldsymbol{\theta})]\right)^2 = \mathrm{Var}_f[\beta(\boldsymbol{\theta})] \geq 0$$

The non-negativity of $\frac{d\phi}{d\lambda}$ follows from the non-negativity of variance. Thus, the lower and upper bounds for ϕ are obtained, respectively, by setting $\lambda = 0$ and $\lambda \to \infty$. For $\lambda = 0$, it is easy to verify that the distribution $f(\boldsymbol{\theta}|P)$ collapses to a uniform distribution. On the other hand, when $\lambda \to \infty$, the distribution concentrates on the largest $\beta(\boldsymbol{\theta})$ value. This is because $e^{\lambda(\beta(\boldsymbol{\theta})-\beta^*)} \to 0$ for any $\beta(\boldsymbol{\theta}) < \beta^* = max_{\boldsymbol{\theta}\in\Theta^P}\beta(\boldsymbol{\theta})$, and therefore $\lim_{\lambda\to\infty}\phi = \max_{\boldsymbol{\theta}\in\Theta^P}\beta(\boldsymbol{\theta})$. QED

Lemma 4 implies that the lower bound of ϕ arises under a uniform distribution over Θ^P (i.e., $\lambda = 0$), while the upper bound corresponds to the maximum value of $\beta(\boldsymbol{\theta})$ over Θ^P. Notably, both bounds are independent of λ. We therefore formulate the following constrained optimization problem:

$$\begin{aligned} &\max_{\sigma\geq 0, \boldsymbol{\phi}} \ \mathcal{L}(\sigma, \{\phi_n\}_n) = \sum_{\substack{m\in Q(x_t,S_t)\\ t=1,\cdots T}} \log\left(\frac{e^{\sigma\phi_m}}{\sum_n e^{\sigma\phi_n}}\right) \\ &s.t. \ \ l_n \leq \phi_n \leq u_n \end{aligned} \tag{9}$$

where l_n and u_n are the lower and upper bounds determined by Lemma 4.

PROPOSITION 1. *If $\sigma^*, \{\phi_n^*\}, \{\lambda_n^*\}$ solve the full likelihood problem* (8)*, then $\sigma^*, \{\phi_n^*\}$ also solve the constrained problem* (9)*. Conversely, if $\sigma^*, \{\phi_n^*\}$ solve* (9)*, then for each ϕ_n^* there exists a unique λ_n satisfying* (7) *that recovers the optimal solution to* (8)*. In other words, $\sigma^*, \{\phi_n^*\}$ are the optimal solutions to the maximum likelihood problem* (8)*.*

This proposition establishes a one-to-one relationship between the solutions to the constrained optimization problem (9) and to the maximum likelihood problem (8) in the $(\sigma, \boldsymbol{\phi})$ space. For each point in this space, there exists a unique $\boldsymbol{\lambda}$ maximizing the likelihood. This motivates the two-stage estimation strategy to obtain the maximum likelihood estimation for $\sigma, \boldsymbol{\phi}, \boldsymbol{\lambda}$ parameters.

THEOREM 3. *The parameters $\sigma^*, \{\lambda_n^*\}$ solve the two-stage procedure in Algorithm* (1) *if and only if they solve the original maximum likelihood problem* (8)*.*

Algorithm 1 The two-stage procedure for estimating $\sigma, \boldsymbol{\phi}, \boldsymbol{\lambda}$ by decomposing the likelihood problem (8)

1: **procedure** TwoStageSolution

Input: $\{(x_t, S_t)\}_{t=1,2,\cdots,T}$

Output: σ, λ_n for every $P_n \in \mathcal{P}$

2: $l_n \leftarrow \mathbb{E}_{uni(\Theta^{P_n})}[\beta(\boldsymbol{\theta})]$ ▷ Expected value w.r.t a uniform distribution over Θ^{P_n}

3: $u_n \leftarrow max_{\boldsymbol{\theta}\in\Theta^{P_n}}\beta(\boldsymbol{\theta})$

4: Solve (9) for σ and ϕ_n, for all $P_n \in \mathcal{P}$

5: σ^* and $\phi_n^* \leftarrow$ solution from the previous step

6: **for all** $P_n \in \mathcal{P}$ **do**

7: Find λ_n by solving $\phi_n^* = \frac{\int_{\Theta^{P_n}} \beta(\boldsymbol{\theta}) e^{\lambda_n \beta(\boldsymbol{\theta})} d\boldsymbol{\theta}}{\int_{\Theta^{P_n}} e^{\lambda_n \beta(\boldsymbol{\theta})} d\boldsymbol{\theta}}$

8: $\lambda_n^* \leftarrow$ solution

9: **end for**

10: **Return** σ^*, λ_n^* for all $P_n \in \mathcal{P}$

11: **end procedure**

This result follows from Lemma 4 and the strict monotonicity of the mapping from λ to ϕ in (7). Bounding ϕ in Step 4 of Algorithm (1) ensures that the feasible regions of both problems coincide. Moreover, the monotonicity of ϕ in λ in (7) guarantees the existence of a unique λ in Step 7.

To implement this approach, all integrals in Algorithm (1) are approximated via Monte Carlo sampling from the parameter space Θ. Because subspaces Θ^P can differ significantly in volume, we ensure uniform representation (i.e., an equal number of samples from each subspace) by identifying the centroid of each Θ^P and initiating a random walk from it with an adaptive step size parameter. The centroid is obtained by solving the following linear program:

$$\begin{aligned} &\max_{\boldsymbol{\theta}\in\Theta^P} \quad \nu \\ &u(x;\boldsymbol{\theta}) - u(y;\boldsymbol{\theta}) \geq \nu,\ \forall u \in \mathcal{U}, \forall (x,y) \in P. \end{aligned} \tag{10}$$

This max-min LP identifies the innermost point of Θ^P, i.e., vector $\boldsymbol{\theta}$ – and hence function $u \in \mathcal{U}$ compatible with P when ordering the elements of X based on their utilities – maximally distant from its defining hyperplanes. [3] We use Algorithm (2) to perform a random walk from the centroid to obtain a desired number of samples from Θ^P:

[3] Our formulation is related to prior work on estimating preference parameters under uniform priors (Srinivasan and Shocker 1973, Jacquet-Lagreze and Siskos 1982, Toubia et al. 2003). As discussed by Toubia et al. (2003), the innermost point minimizes some expected error in estimating preference parameters if we assume a uniform distribution over the parameter space.

Algorithm 2 Sampling from Θ^P subspaces.

1: **procedure** UNIFORMSAMPLING

Input: k_s: the number of samples from each subspace $\Theta^P \subset \Theta$, δ_0: the sampling initial step size, $\alpha \in (0,1)$: the adaptive step size parameter, ϵ: a small number determining the smallest admissible step size.

Output: $Sample_P = \{\boldsymbol{\theta}_k\}_{k=1,2,\cdots,k_s}$ where $\boldsymbol{\theta}_k \in \Theta^P$, for each $P \in \mathcal{P}$

2: **for all** $P_n \in \mathcal{P}$ **do**

3: $k \leftarrow 0$, $reject \leftarrow 0$, $\delta = \delta_0$

4: Solve (10) to obtain $\boldsymbol{\theta}^*$ and v^* ▷ Finding centroid of Θ^P

5: **if** $v^* > 0$ **then** $\boldsymbol{\theta}_1 \leftarrow \theta^*$, $Sample_P \leftarrow \boldsymbol{\theta}_1$, and $k \leftarrow k+1$

6: **else** Eliminate P_n from $\mathcal{P}$ and go to step 2

7: **end if**

8: **while** $k < k_s$ and $reject < 100$ **do**

9: Compute $\boldsymbol{\theta} = \boldsymbol{\theta}_k + \delta \boldsymbol{r}$ where $\boldsymbol{r}$ is a random unit vector

10: **if** $\boldsymbol{\theta} \in \Theta^P$ **then** $\boldsymbol{\theta}_{k+1} \leftarrow \boldsymbol{\theta}$, $Sample_{P_n} \leftarrow Sample_{P_n} \cup \boldsymbol{\theta}_{k+1}$, $k \leftarrow k+1$, $reject \leftarrow 0$

11: **else** $reject \leftarrow reject + 1$

12: **if** $reject > 100$ **then** $\delta = \alpha\delta$

13: **if** $\delta < \epsilon$ **then** $\delta = \delta_0$, $reject \leftarrow 0$

14: **end if**

15: **end if**

16: Go back to step 9

17: **end if**

18: **end while**

19: **end for**

20: **end procedure**

The procedure described in Algorithm (2) initiates a random walk from the centroid of each subspace Θ^P and collects a predetermined number of samples, denoted by k_s. Given that the volumes of different Θ^P subspaces can vary substantially, Steps 17 to 22 in the algorithm dynamically adjust the step size δ to ensure efficient exploration. Specifically, if the number of consecutively rejected samples – i.e., proposals that fall outside Θ^P – exceeds a threshold (set to 100), the step size δ is

reduced by a factor α (e.g., 0.75). To avoid oversampling a narrow region of Θ^P, the algorithm resets δ to its initial value δ_0 (set to 0.05) when it becomes too small.

A key advantage of this approach is that it ensures an equal number of samples from each Θ^P, thereby guaranteeing equal representation of each ranking P in the numerical procedure described in Algorithm (1). Furthermore, because the sampling process constitutes a Markov chain, the random walk over Θ^P converges to a uniform distribution over the subspace as k_s increases. This property is essential, as uniform sampling over the unpartitioned space Θ would not yield the same representational balance.

To further reduce computational cost and sample autocorrelation while maintaining accuracy, our implementation utilizes a thinning strategy (Brooks et al. 2011), by selecting one point from every batch of 10 sampled points. Using the outputs from Algorithm (2), Algorithm (1) can be employed to estimate the parameters σ and λ_n, as well as the full distribution f over Θ and the induced distribution $\mu \in \Delta(\mathcal{P})$.

The sampling process over the partitioned space $\times$ is the most computationally intensive component of our estimation method. However, it is performed only once. Once the sample set $\{\boldsymbol{\theta}\}$ is collected and stored, future computations – including model updates in response to new data or different preference structures P – require only recomputation of $\boldsymbol{\phi}$ via Eq. (7).

Before presenting the numerical results, we describe a distance-based estimation procedure as an alternative to the maximum likelihood method. Unlike the likelihood-based approach, which requires choice observations at the choice task level, the distance-based method can be applied even when such data are available only at the aggregate level as a choice share of options in each menu.

4.1. Distance-Based Estimation

Recall that we consider using data for model estimation that consists of pairs (x_t, S_t) that represent the choice of x_t from menu S_t. Define the matrix A such that each element a_{tn} equals 1 if the ranking P_n is compatible with choosing x_t from the menu S_t – that is, if $x_t P_n y$ for all $y \in S_t \setminus \{x_t\}$. Let the column vector $\boldsymbol{\mu}$ denote a probability distribution over the set of rankings, where the n-th element μ_n equals $\mu(P_n)$. Then, the inner product of the t-th row of A, denoted by $\boldsymbol{a}_t$, and the vector $\boldsymbol{\mu}$ yields the predicted choice probability of x_t from S_t, i.e., $\boldsymbol{a}_t \boldsymbol{\mu} = \rho_\mu(x_t, S_t)$ by construction.

The objective of the distance-based estimation approach is to identify a probability vector $\boldsymbol{\mu}$ that produces predicted choice probabilities ρ_μ as close as possible to the observed choice shares ρ. This is achieved by solving the following constrained optimization problem:

$$\begin{aligned} \min_{\boldsymbol{\mu}\geq 0,\sigma\geq 0,\boldsymbol{\phi},z} \quad & d(\rho_\mu,\rho) = \|z-\boldsymbol{\rho}\|_2^2 \\ s.t. \quad & A\boldsymbol{\mu} - z = \mathbf{0} \\ & \boldsymbol{\mu} - \frac{e^{\sigma\boldsymbol{\phi}}}{\mathbf{1}^\tau e^{\sigma\boldsymbol{\phi}}} = \mathbf{0} \\ & \boldsymbol{\phi} - \boldsymbol{u} \leq \mathbf{0} \\ & \boldsymbol{l} - \boldsymbol{\phi} \leq \mathbf{0} \end{aligned} \tag{11}$$

where z denotes the vector of choice probabilities induced by μ, $\boldsymbol{\rho}$ represents the vector of observed choice shares, and $\boldsymbol{l}$ and $\boldsymbol{u}$ are the lower and upper bounds on $\boldsymbol{\phi}$ as defined in Lemma 4. The first constraint ensures that z corresponds to the choice probabilities generated by ρ_μ, and the next three constraints impose the RPM structure described in Section 3.2.

Note that the matrix A encapsulates all the available information from the observed choice data, as it maps observed choices to the space of rankings. For this reason, we refer to A as the *information matrix*. This matrix is typically high-dimensional, with the number of columns growing factorially with the size of X, i.e., $O(|X|!)$, which renders the estimation problem computationally challenging.

To address this challenge, we introduce two approximation strategies that reduce computational complexity at the cost of estimation accuracy. The first approach restricts the support of the RPM by constructing μ over a subset of $\mathcal{P}$ rather than the full permutation set. The other approach limits each ranking to its top-k elements, effectively working with partial permutations.

4.2. Restricted-Support RPM

Let us refine our notation by defining $\mathcal{P}_0$ as the set of all complete, transitive, and asymmetric binary relations on X, and let $\mathcal{P} \subseteq \mathcal{P}_0$. When $\mathcal{P}$ is a proper subset of $\mathcal{P}_0$, the pair $(\mu, \mathcal{P})$ constitutes an RPM with restricted support. By controlling the size of $\mathcal{P}$, one can effectively manage the computational complexity of the optimization problems (9) and (11). Notably, both these optimization formulations and their corresponding solution methods – Algorithms (1) and (2) — can be applied without modification under this restriction.

We consider two approaches for constructing $\mathcal{P}$. The first involves random sampling from $\mathcal{P}_0$, where the sample size mediates the trade-off between computational efficiency and approximation accuracy. This procedure is equivalent to sampling from the column space of the information matrix A as defined in Section 4.1. Importantly, the sampling process begins by partitioning the

preference space such that each partition consists of rankings where a specific alternative occupies the top rank. One ranking is initially sampled from each partition, after which additional rankings are sampled uniformly at random from $\mathcal{P}_0$ until the desired sample size is reached. This adjustment ensures that $\rho(x, S)$ is not inadvertently set to zero, which would otherwise result in a degenerate likelihood function.

The second approach constructs $\mathcal{P}$ by identifying a basis for the column space of the matrix A. The rationale is that since the vector of predicted choice shares $\rho_\mu \equiv A\boldsymbol{\mu}$ is a linear combination of the columns of A, then, if a column $\boldsymbol{a}_n$ can be written as a convex combination of other columns, i.e., $\boldsymbol{a}_n = \sum_{j=1}^{J} t_j \boldsymbol{a}_j$ with $\sum_j t_j = 1$, its contribution to ρ_μ is likewise a linear combination of their respective contributions with the same weighting factors $t_1, t_2, \cdots, t_J$. In this case, if $P_n \in Q(x, S)$, removing P_n from $\mathcal{P}$ redistributes its probability mass μ_n across the basis rankings according to the weights t_j, $j = 1, \cdots, J$. Specifically, for a pair (x, S) and its corresponding row vector $\boldsymbol{a}$ in A, we have $\rho_\mu(x, S) = \sum_{k=1}^{|\mathcal{P}|} a_k \mu_k$. Since $a_n = \sum_j t_j a_j$, each coefficient a_j will effectively receive an increment of $t_j \mu_n$, thereby reallocating a share $t_j \mu_n$ of the removed probability mass to μ_j. As the resulting matrix A constructed from the restricted support $\mathcal{P}$ is full-rank, the corresponding RPM $(\mu, \mathcal{P})$ remains identifiable.

Nonetheless, a limitation of this approach is that the predicted choice shares ρ_μ cannot always be fully recovered when $\mathcal{P}$ is restricted to rankings associated with a linearly independent subset of columns of A. This occurs because the probability mass assigned to a ranking P_j, corresponding to a column a_j in the basis, may implicitly contain contributions from dependent rankings that are not necessarily compatible with the choice situation (x, S). Despite this limitation, the method offers a compelling approximation strategy, as it reduces problem size by eliminating redundant information. In particular, it reduces the number of columns in A from $O(|X|!)$ to $O(T)$, where T is the length of the vector of observed choice shares $\boldsymbol{\rho}$. Since typically $T \ll |X|!$, and in the worst case $T = O(2^{|X|})$, this reduction can yield substantial computational gains.

Finally, to construct the restricted support RPM via the column basis method, it is not necessary to compute or store the full information matrix A, which would be infeasible for large $|\mathcal{P}|$. Instead, candidate rankings are evaluated sequentially. For each candidate, the corresponding column of A is computed and appended to the current matrix if and only if it increases the rank. In this way, the support set $\mathcal{P}$ is incrementally built to ensure linear independence among the associated columns.

4.3. Partial-Permutation RPM

The approximation approach described in this section modifies the elements of $\mathcal{P}$ such that they no longer represent complete, transitive, and asymmetric binary relations on X. Instead, the elements are partial permutations, obtained by restricting the depth of each ranking. Consequently, the size of $\mathcal{P}$ is altered. When only the top-k alternatives in each ranking are considered, the cardinality of $\mathcal{P}$ becomes $|\mathcal{P}| = |X| \times (|X| - 1) \times \cdots \times (|X| - k + 1)$. The underlying intuition is that, when facing a choice problem, the DMs are indifferent among all remaining options if none of their top k preferences are available. This idea bears conceptual resemblance to the notion of a rank cut-off discussed in Bai et al. (2024).

When $\mathcal{P}$ is composed of partial permutations, the matrix A is no longer binary. This is because, for a pair of alternatives $x, y \in S$, a partial permutation may be indeterminate in revealing a strict preference relation. As a result, it is not always immediately clear whether a given partial ranking P is compatible with an observed choice instance (x, S). To determine the value of the entry a_n in the row vector $\boldsymbol{a}$ of matrix A corresponding to the pair (x, S), we consider the following three cases:

- If x is present in P_n, and all elements of $S \setminus \{x\}$ are either ranked below x or do not appear in P_n, then $a_n = 1$.
- If x is not present in P_n and at least one element of $S \setminus \{x\}$ is present, or if x appears in P_n and at least one element of $S \setminus \{x\}$ is ranked above x, then $a_n = 0$.
- If none of the elements in S are present in P_n, then $a_n = 1/|S|$.

The final case reflects a situation in which the customer type represented by P_n chooses uniformly at random from the set S, due to none of the alternatives in S being among her top k ranked options.

5. Prototype Analysis

Consider a dataset consisting of $T = 100$ binary choices among $|X| = 5$ alternatives, each described by $M = 4$ attributes, with each attribute taking one of $L = 4$ possible levels. The choice data are generated using randomly constructed options and simulated utility functions. Each option is represented by a vector $\boldsymbol{x}_j = (x_{1j}, \ldots, x_{4j})$, where each component denotes the level of the corresponding attribute. Utilities are computed using a simulated utility function of the form

$$U_j = \sum_{m=1}^{4} w_m u_{mj} + \epsilon,$$

which comprises two components: a weighted sum of marginal utilities and a stochastic noise term. The marginal utility functions u_{mj} are independently generated and scaled to lie within the

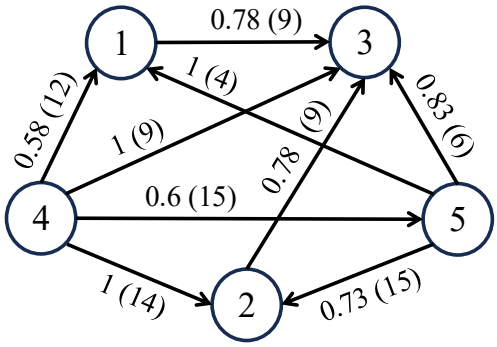

Figure 3 **Summary of the generated choice data in the illustrative example. Circles represent options, and for each binary menu $\{x, y\}$, the label over the arrow from x to y includes the fraction of times x is chosen from the menu, and the number in parentheses shows the number of times the menu is observed.**

interval $[0, 1]$. The weights $\boldsymbol{w} = (w_1, \ldots, w_4)$ are drawn from a random distribution such that $w_m \in (0, 1)$ and $\sum_{m=1}^{4} w_m = 1$, ensuring a convex combination of the marginal utilities. Here, u_{mj} represents the marginal utility of option j with respect to attribute m, based on its assigned level and the corresponding randomly generated utility function. The stochastic component $\epsilon \sim \mathcal{N}(0, s^2)$ introduces randomness into the utility, with a mean of 0 and standard deviation s. Given that utilities are normalized within unit intervals, we set $s = 0.2$, corresponding to a moderate to high noise level. In our illustrative example, this noise level led to approximately 24% of observed choices being inconsistent with the simulated underlying utility values. Descriptions of the options, simulated weight vectors, and resulting utility values are presented in Table 2. A summary of the generated choice data is provided in Figure 3.

Table 2 **Descriptions of the simulated options, simulated weights, and ground truth utility values in the illustrative example.**

	attribute 1	attribute 2	attribute 3	attribute 4	utilities
weights	0.4648	0.2991	0.0542	0.1818	
option 1	0	3	3	2	0.3828
option 2	1	1	1	2	0.3221
option 3	2	0	1	3	0.2761
option 4	2	2	0	3	0.5320
option 5	3	0	0	2	0.4942

We first estimate the RPM $(\mu, \mathcal{P}_0)$ with full support using both the maximum-likelihood and distance-based methods. We then present the results for the restricted support RPM for both settings

Table 3 Performance metrics for all seven RPM variants, including attribute weights and their ranges, mean absolute error of the attribute weight estimates, and the log likelihood. Results for the restricted support with random sampling include means and standard deviations across 20 independent replications.

	attribute 1	attribute 2	attribute 3	attribute 4	MAE	Log-likelihood
True weights	0.465	0.299	0.054	0.182		
Full support	0.455 (0.242, 0.589)	0.315 (0.079, 0.411)	0.084 (0.012, 0.308)	0.146 (0.024, 0.522)	0.023	-44.7
Restricted support random with $\|\mathcal{P}\| = \|\mathcal{P}_0\|/3$	0.429 (0.211, 0.567)	0.311 (0.082, 0.430)	0.099 (0.017, 0.333)	0.162 (0.023, 0.540)	0.028	-45.8 (std = 0.96)
Restricted support information matrix basis (MLE)	0.427 (0.208, 560)	0.263 (0.091, 0.440)	0.125 (0.024, 0.325)	0.186 (0.019, 0.538)	0.037	-52.3
Restricted support information matrix basis (min distance)	0.455 (0.285, 0.710)	0.315 (0.042, 0.313)	0.084 (0.018, 0.310)	0.146 (0.020, 0.516)	0.023	-52.6
Top-k (k=2)	0.407 (0.152, 0.649)	0.305 (0.041, 0.528)	0.119 (0.011, 0.407)	0.169 (0.011, 0.605)	0.035	-44.8
Restricted-support random and top-k (k=2, $\|\mathcal{P}\| = 10$)	0.405 (0.157, 0.630)	0.269 (0.037, 0.500)	0.133 (0.013, 0.396)	0.192 (0.012, 0.597)	0.045	-53.1 (std = 4.03)
Restricted-support information matrix basis and top-k (k=2)	0.399 (0.153, 0.649)	0.292 (0.035, 0.528)	0.132 (0.011, 0.428)	0.178 (0.010, 0.599)	0.039	-48.8

Table 4 Performance metrics for all seven RPM variants, including estimated choice probabilities across all menus, choice probabilities conditioned on menus consistent/inconsistent with ground-truth preferences, and recovered preference rankings. Results for the restricted support with random sampling include means and standard deviations across 20 independent replications.

	Market share recovery			Underlying preference recovery	
	all menus	consistent choices	inconsistent choices	$\mathbb{E}_\mu\big(\tau(P, P^*)\big)$	$\tau(P^{rep}, P^*)$
Full support	0.703	0.818	0.340	0.583	1.000
Restricted support random with $\|\mathcal{P}\| = \|\mathcal{P}_0\|/3$	0.686 (std = 0.009)	0.791 (std = 0.017)	0.352 (std = 0.023)	0.497 (std = 0.063)	0.840 (std = 0.179)
Restricted support information matrix basis (MLE)	0.664	0.780	0.296	0.494	0.800
Restricted support information matrix basis (min distance)	0.674	0.794	0.292	0.517	0.800
Top-k (k=2)	0.664	0.762	0.354	0.446	1.000
Restricted-support random and top-k (k=2, $\|\mathcal{P}\| = 10$)	0.623 (std = 0.033)	0.706 (std = 0.056)	0.378 (std = 0.051)	-0.188 (std = 0.233)	-0.327 (std = 0.413)
Restricted-support information matrix basis and top-k (k=2)	0.670	0.776	0.333	-0.301	-0.600*

where $\mathcal{P} \subset \mathcal{P}_0$ is constructed via random sampling from the preference space $\mathcal{P}_0$, and via identifying the basis for column space of the information matrix A. In both cases, we estimate the restricted-support RPM using the maximum-likelihood and distance-based methods. Next, we present the results for the partial-permutation RPM with $k = 3$, that is, when only top-3 options in each ranking are considered. Finally, combining the two approximation strategies, we present the results for partial-permutation RPM with restricted support based on random sampling and the information matrix column basis. Summary statistics for all seven RPM variants are provided in Tables 3 and 4.

We begin by presenting the results of our stylized example for the RPM with full support, estimated via maximum likelihood. The average estimated choice probability across the $T = 100$ observed (menu, choice) pairs is 0.703; higher values indicate a better fit to the data. Some observed

choices are inconsistent with the ground-truth simulated preferences due to the inclusion of a stochastic error component. The average estimated choice probability among consistent choices is 0.818, while for inconsistent choices it is only 0.340, highlighting the RPM's ability to accommodate noise in observed behavior.

To evaluate how well the estimated RPM recovers the underlying preference structure, we compared the estimated distribution over the set of rankings with the ground-truth ranking P^*, which is consistent with the true utility values of the alternatives. We assessed this comparison using two metrics. First, we computed the expected Kendall's tau correlation between the estimated RPM and the true ranking, i.e.:

$$\mathbb{E}_\mu\Big(\tau(P, P^*)\Big) = \sum_n \mu(P_n)\tau(P_n, P^*), \tag{12}$$

and obtained an expected correlation of 0.583 (p-value = 0.117) for the estimated RPM. This implies that a ranking randomly drawn from $\mathcal{P}$, according to the distribution μ, matches the true underlying ranking with an expected Kendall's tau correlation of 0.583 under the noisy conditions of our illustrative example.[4]

Second, in settings involving heterogeneous agents with potentially conflicting preferences – such as those studied in social choice theory – it is often desirable to identify a representative ranking that best aggregates the preferences of the population. To this end, we adopt the assent-maximizing welfare function introduced by Baldiga and Green (2013), which generalizes the well-known Kemeny rule:

$$P^{rep} = \arg\max_P \sum_n \mu(P_n)\tau(P, P_n). \tag{13}$$

We then compare the representative preference P^{rep} with the ground-truth ranking P^* using Kendall's tau correlation. In our illustrative example, we obtain a perfect correlation of 1.0 (p-value = 0.006) [5], indicating that the assent-maximizing ranking is identical to the true underlying ranking presented in the last column of Table 2. It is worth noting that the representative preference derived from Eq. (13) may fall outside the RPM support in the case of restricted support RPM.

Moreover, the RPM framework enables the construction of a probability distribution over each component of $\boldsymbol{\theta}$ – the partworths – as well as over any function thereof, such as attribute weights. This

[4] Note that Kendall's tau ranges from −1 to 1.

[5] To compute p-values for the Kendall's tau correlations, we used an exact permutation test, which is feasible due to the small number of possible rankings ($N!$). We enumerated all possible rankings and computed their Kendall's tau correlation with the ground truth ranking. The p-value was then defined as the fraction of rankings whose correlation coefficient with the ground-truth ranking exceeds the observed correlation coefficient.

is accomplished by evaluating the conditional probability, as defined in Eq. (6), for each sampled point $\boldsymbol{\theta} \in \Theta^P$, where Θ^P is a preference subspace obtained using Algorithm (2). Each sample is weighted according to the estimated probability mass $\mu(P)$. Specifically, for any component θ_j of the utility parameter vector $\boldsymbol{\theta}$ and a given quantile q:

$$\mathbb{P}(\theta_j \leq \boldsymbol{q}) = \sum_n \hat{\mu}(P_n) \frac{\sum_{\boldsymbol{\theta}_k \in P_n : \theta_{jk} \leq q} e^{\hat{\lambda}_n \beta(\boldsymbol{\theta}_k)}}{\sum_{\boldsymbol{\theta}_k \in P_n} e^{\hat{\lambda}_n \beta(\boldsymbol{\theta}_k)}}. \tag{14}$$

Note that Algorithm (2) is specifically designed to approximate uniform sampling from the subspaces of Θ^P. As a result, Eq. (14) provides an approximation of the cumulative distribution function obtained by combining the conditional probability in Eq. (6) with the probability in Eq. (5). The accuracy of this approximation improves with larger sample sizes k_s in the sampling algorithm. Since the inferred probabilities over the parameters are derived from sampled points in the parameter space, it is also possible to compute bounds on the corresponding cumulative distribution functions. This is done by identifying the smallest and largest sampled values of each parameter within a subspace and determining whether a given quantile falls below, within, or above this range. Consequently, we can characterize the entire range of possible cumulative distributions for each parameter (i.e., partworth) or for any function of these parameters (e.g., attribute weights). Figure 4 illustrates the cumulative probability distributions and their corresponding lower and upper bounds for the attribute weights in our illustrative example.

The estimated probability distributions enable the computation of expected values, along with their corresponding lower and upper bounds, for parameters of interest. These bounds are obtained by evaluating the upper and lower cumulative probability distributions, respectively. Figure 5 presents the true attribute weights (i.e., the ground truth derived from the simulated data), their estimated expected values under the full-support RPM, and the associated lower and upper bounds. The mean absolute difference between the estimated expected weights and their true values is 0.023, demonstrating the RPM's effectiveness in recovering attribute weights in our illustrative example. As a benchmark, we computed the mean absolute difference for the MNL estimates, which yielded a value of 0.255, nearly an order of magnitude larger than that of the RPM, even though MNL correctly recovered the ground truth ranking. Furthermore, the mean absolute difference between the RPM-estimated market shares and the observed market shares – shown in Figure 1 – is 0.036, with individual differences ranging from 0 to 0.075. The results for the distance-based estimation were nearly identical; thus, we omit their presentation for brevity.

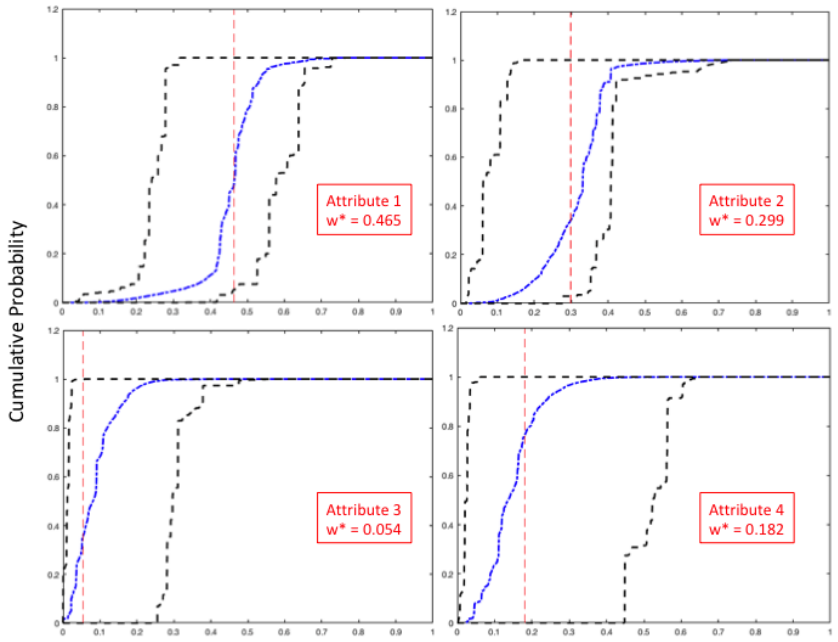

Figure 4 **(Color online) Cumulative probability distributions and their minimum and maximum bounds for the attribute weights in the illustrative example. Vertical red dashed lines indicate the true attribute weights.**

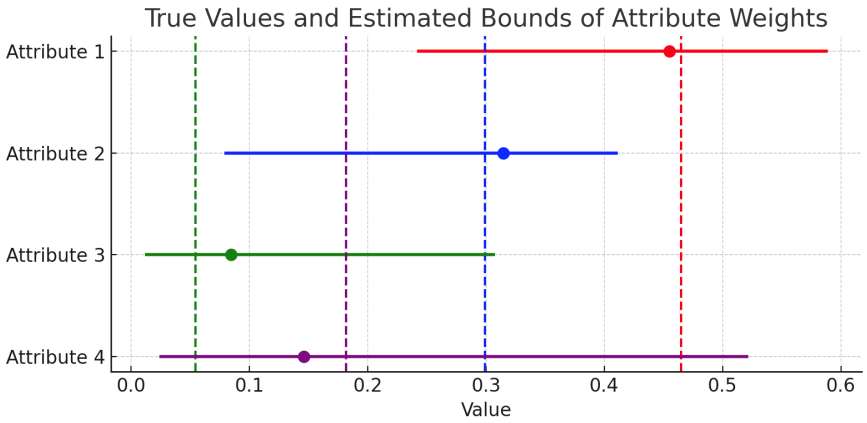

Figure 5 **(Color online) The expected value, minimum, and maximum bounds of attribute weights for maximum likelihood estimated RPM with full support. Vertical dashed lines indicate the true attribute weights.**

We next present the results of our illustrative example for the restricted-support RPM, using two approaches for constructing the support set $\mathcal{P}$: (i) random sampling from the full permutation space $\mathcal{P}_0$, and (ii) identifying a basis for the column space of the information matrix A. In the random sampling approach, we select a restricted support of size $|\mathcal{P}| = |\mathcal{P}_0|/3$ from the full preference space $\mathcal{P}_0$. To mitigate potential dependence on a specific sample configuration, we repeat this procedure

20 times. All results reported for the restricted-support RPM with random sampling are averaged over these 20 repetitions.

We observe that the log-likelihood decreases from −44.7 in the full-support RPM to −45.8 (standard deviation = 0.96) in the restricted-support RPM with random sampling. This decline is expected, given that the full-support RPM subsumes the restricted version and involves a larger parameter space. Nonetheless, the relatively small difference in log-likelihood, despite a substantially reduced support size, is notable and suggests that restricting the support can be an effective means of controlling model complexity. The average estimated choice probability is 0.686 (standard deviation = 0.009). For observed choices consistent with the ground truth, this value is 0.791 (standard deviation = 0.017), and for inconsistent choices, it is 0.352 (standard deviation = 0.023). While these figures indicate a modest reduction in performance relative to the full-support RPM, the differences are negligible.

When constructing $\mathcal{P}$ based on the basis of the column space of the information matrix, the support size is reduced to $|\mathcal{P}| = 11$, which is nearly four times smaller than the random sampling case and nearly eleven times smaller than the full-support RPM. Maximum likelihood estimation for this restricted-support RPM yields a log-likelihood of −52.3, which is lower than in the previous cases. However, the model still performs well in terms of attribute weight estimation, market share recovery, and preference recovery. Furthermore, in this configuration, the distance-based estimation method yields tighter bounds on the attribute weights compared to the maximum likelihood approach.

Notably, the partial permutation approximation with $k = 2$ significantly simplifies the model while achieving performance comparable to the restricted-support RPM, except in estimating the bounds on attribute weights. The results for the partial permutation RPM indicate greater inferential uncertainty, as reflected in wider bounds on the attribute weights.

Finally, combining the two approximation strategies – restricted support and partial permutations – did not substantially degrade performance. For example, this configuration yielded a mean attribute weight estimation error of 0.039 and a log-likelihood of −48.8, compared to 0.023 and −44.7, respectively, for the full-support RPM. However, this was the only setting in which the estimated representative preferences were not positively correlated with the true underlying preference.

For comparison, we also estimated an MNL model using the same data. The MNL yielded a log-likelihood of -50.2 and a mean absolute error of 0.255 for the estimated weights—substantially larger than the error for any RPM variant[6], and correctly recovered the underlying true ranking.

6. Empirical Analysis

For our empirical analysis, we used the dataset from Rooderkerk et al. (2011), derived from a carefully designed choice-based conjoint experiment. In each choice task, participants selected one option from a menu comprising three hypothetical digital cameras, each characterized by two attributes: picture quality (measured in megapixels) and optical zoom. Both attributes varied across five levels. The experiment was constructed to ensure the absence of dominated alternatives in any menu. Additionally, some menus were intentionally designed to include asymmetrically dominated options, aimed at eliciting context-dependent preferences that enhance the attractiveness of a target option. The resulting dataset contains $1,915$ choice observations across 30 unique menus, involving a total of 18 distinct digital cameras.

We randomly partitioned the 30 menus into training and test sets. The training set included only the choice data from menus assigned to the training condition, while the test set comprises the remaining menus. We considered two training set sizes: a smaller set consisting of 12 menus (40% of the data) and a larger set containing 24 menus (80%).

For each training set, we first estimated an MNL model to serve as a benchmark and then derived models from the various RPM specifications. Specifically, we considered two approaches for constructing the support of the RPM. The first involves random sampling of partial rankings from the set of all possible partial permutations, with the number of sampled rankings varied across conditions. The other approach identifies a set of partial rankings whose corresponding columns span the column space of the information matrix; we refer to this as the *orthogonal basis method*. For random sampling, we employed a Markov chain sampling procedure that iteratively swaps the positions of two alternatives in a ranking, using a step size of 10 to promote diversity among the sampled rankings. In both methods, we ensured that no dominated alternative was ranked above its dominant counterpart.

[6] This is more than five times the largest MAE among the seven RPM variants. The MNL estimated weights are $w = (0.080, 0.419, 0.444, 0.057)$. Despite the fact that the data-generating process was a probit-based random utility model with i.i.d. errors – closely resembling logit – the MNL model did not outperform the RPM in market share recovery in this particulare xample with limited data. Specifically, the estimated choice probabilities were 0.670, 0.790, and 0.286 for all menus, consistent choices, and inconsistent choices, respectively. These results suggest the RPM's ability to recover market share patterns accurately, even when the data are generated from a distinct stochastic process.

We fixed the ranking depth at $k = 5$ and applied the distance-based estimation method. For this depth, the orthogonal basis method yields 19 distinct partial rankings. Under the random sampling approach, we experimented with support sizes of 10, 20, and 30 sampled rankings. To assess the effect of ranking depth, we repeated the orthogonal basis analysis with $k = 7$ under both training set sizes. This configuration produces 30 partial rankings – matching the largest support size used in the random sampling condition. For comparability, we also perform random sampling with 30 rankings at $k = 7$. For each configuration, we repeated the procedure 20 times, drawing random splits of the 30 distinct menus into training and test sets, and estimating both RPM and MNL models on identical data partitions.

We evaluated model performance separately on the training and test sets using two key metrics, (i) the mean absolute error (MAE) between predicted and observed choice probabilities across menus, and (ii) in-sample and out-of-sample (holdout) hit rates. The hit rate is defined as the proportion of choice tasks in which the alternative with the highest predicted choice probability matches the observed choice. In cases of ties in predicted probabilities, we break the tie using the Jaccard similarity index, equivalent to random selection among tied alternatives. MAE is calculated over the 30 menus (divided between training and test sets), capturing each model's accuracy in predicting market shares. In contrast, hit rates are computed over all $1,915$ individual choice observations and therefore account for the frequency with which menus appear in the dataset. Our resulted are presented in Table 5.

Table 5 Mean RPM Hit Rates and MAE Values by RPM Type and Training Size.

Train Size	RPM Type	In-sample		Holdout	
		Hit Rate (%)	MAE	Hit Rate (%)	MAE
Small	orthogonal basis and top5	51.9	0.194	50.5	0.213
	orthogonal basis and top7	**65.4**	0.182	**61.9**	0.224
	random with $\|\mathcal{P}\| = 10$ and top5	57.1	0.184	46.2	0.227
	random with $\|\mathcal{P}\| = 20$ and top5	64.6	0.170	54.4	0.210
	random with $\|\mathcal{P}\| = 30$ and top5	63.7	0.173	48.1	**0.208**
	random with $\|\mathcal{P}\| = 30$ and top7	64.9	**0.163**	55.6	0.210
	MNL	44.7	0.350	41.5	0.357
Large	orthogonal basis and top5	55.8	0.198	49.0	0.202
	orthogonal basis and top7	**66.5**	0.200	**63.3**	0.204
	random with $\|\mathcal{P}\| = 10$ and top5	49.9	0.204	51.7	0.211
	random with $\|\mathcal{P}\| = 20$ and top5	59.0	0.185	56.9	0.196
	random with $\|\mathcal{P}\| = 30$ and top5	65.6	**0.171**	58.1	0.197
	random with $\|\mathcal{P}\| = 30$ and top7	62.6	0.177	62.8	**0.185**
	MNL	43.6	0.349	44.3	0.346

RPM methods exhibit significantly superior performance compared to the benchmark MNL model in both in-sample and holdout evaluations. Across both small and large training set conditions,

RPM achieves hit rates that are approximately 20% higher than those of the MNL (in absolute terms). In terms of predictive accuracy, as measured by MAE, RPM models reduce error by nearly 50% relative to MNL in both evaluation settings, indicating a substantial improvement.

Among the RPM variants, the orthogonal basis method consistently yields the highest prediction accuracy, achieving the best in-sample and holdout hit rates across both training set sizes. This advantage is particularly pronounced in the small training set condition, where this RPM variant outperforms the random sampling approach by a more noticeable margin. When comparing methods with identical support size and ranking depth ($k = 7$), the difference in holdout hit rate between the basis method and the random sampling method with 30 rankings is negligible under the large training set condition (0.5% difference; p-value = 0.922), but becomes more substantial under the small training set condition (exceeding 6%; p-value = 0.051). Figure 6 summarizes the hit rate differences between the RPM variants and the MNL benchmark across training set sizes and evaluation conditions.

Regarding the accuracy of estimated choice probabilities, ρ, the random sampling method with 30 rankings of depth $k = 5$ or $k = 7$ yields the lowest MAE. This method outperforms the orthogonal basis approach in terms of MAE, particularly when the training set is small. However, as the training set size increases, the performance gap between the two RPM variants diminishes, with both methods exhibiting comparable accuracy. These trends are illustrated in Figure 7.

An important exogenous parameter in the RPM top-k method is the ranking depth k. To evaluate how the random sampling and orthogonal basis methods respond to variations in this parameter, we compared the two RPM variants incorporating the orthogonal basis and random sampling methods. When $k = 5$, the orthogonal basis method yields 19 distinct rankings, whereas increasing the depth to $k = 7$ expands the set to 30 rankings. For a consistent comparison, we selected the random sampling method with 30 rankings, which previously demonstrated the strongest performance among the random sampling variants.

The results, depicted in Figure 8, indicate that increasing the ranking depth k generally improves hit rate performance for both RPM variants.[7] However, the performance gains are more pronounced for the orthogonal basis method, as evidenced by the steeper improvement in hit rates. Regarding MAE, both methods exhibit similar sensitivity to k for in-sample predictions. In contrast, for holdout MAE, increasing the ranking depth yields greater benefits for the random sampling method, suggesting that deeper rankings enhance generalization performance in this setting.

[7] An exception occurs in the in-sample hit rate for the random sampling method under the large training condition, where a larger k slightly reduces performance.

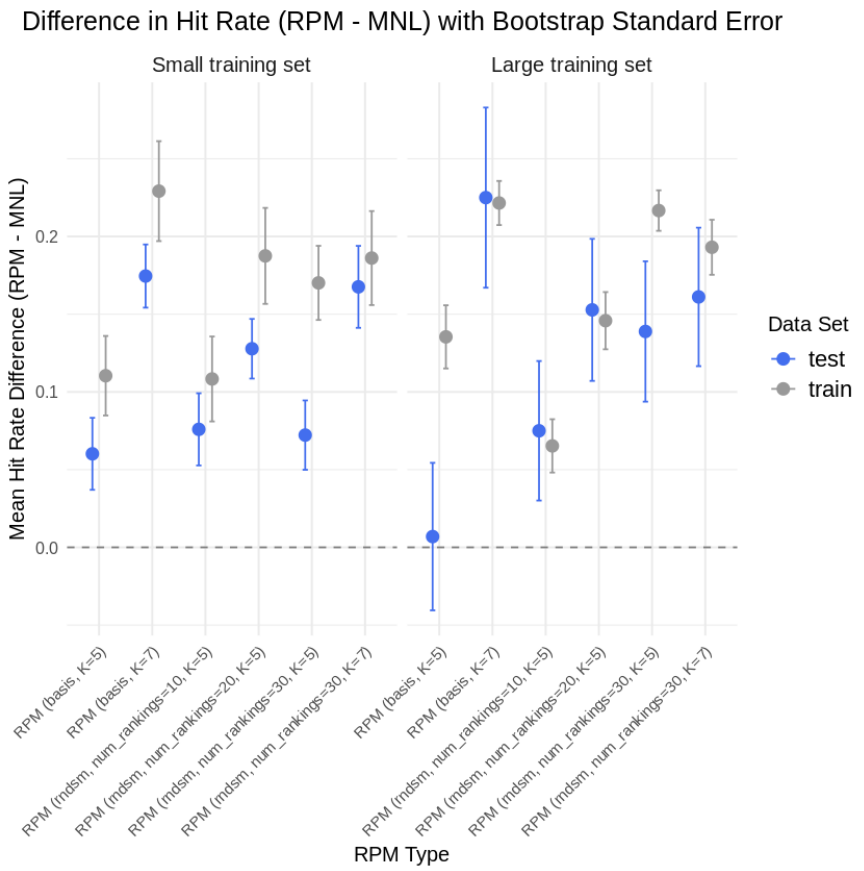

Figure 6 **(Color online) Mean hit rate difference between RPM and MNL in training (gray) and test (blue) sets for small (left) and large (right) training sets, and for different RPM types. The vertical lines represent bootstrap standard errors with** 1000 **replications.**

Finally, we report the RPM computation times separately for Monte Carlo sampling in Algorithm (2) and the solution method in Algorithm (1), where step 4 solves the constrained optimization problem (11), as well as the computation time for MNL. Our results, presented in Table 6, show that the solution time for MNL is considerably smaller than that of RPM when the training set size is small, but there is no noticeable difference for a large training set size. Moreover, the solution time for RPM increases less sharply with the training set size compared to that of MNL. For RPM, both the time required to solve the constrained optimization problem and the time for Monte Carlo sampling increase with the ranking depth k and with the RPM support size. However, the computational bottleneck in RPM estimation lies in sampling from the parameter subspaces.

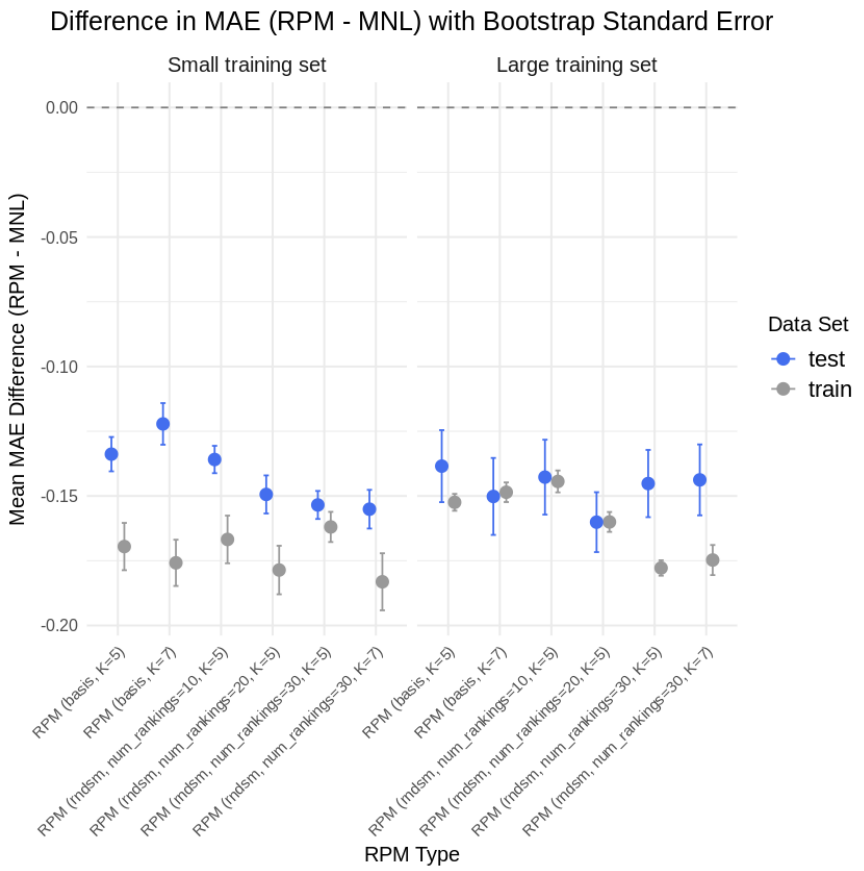

Figure 7 **(Color online) Mean absolute error (MAE) difference between RPM and MNL in training (gray) and test (blue) sets for small (left) and large (right) training sets, and for different RPM types. The vertical lines represent bootstrap standard errors with** 1000 **replications.**

Table 6 **Average Computation Time (in seconds) by RPM Type and Training Size.**

Train Size	RPM Type	RPM Sampling	Solve Time
Small	orthogonal basis and top5	297.706	3.499
	orthogonal basis and top7	541.685	7.561
	random with $\|\mathcal{P}\| = 10$ and top5	164.067	1.775
	random with $\|\mathcal{P}\| = 20$ and top5	348.961	4.484
	random with $\|\mathcal{P}\| = 30$ and top5	521.106	8.202
	random with $\|\mathcal{P}\| = 30$ and top7	482.192	6.127
	MNL		1.998
Large	orthogonal basis and top5	602.714	3.798
	orthogonal basis and top7	1011.348	6.986
	random with $\|\mathcal{P}\| = 10$ and top5	306.710	1.491
	random with $\|\mathcal{P}\| = 20$ and top5	677.872	4.726
	random with $\|\mathcal{P}\| = 30$ and top5	922.781	6.919
	random with $\|\mathcal{P}\| = 30$ and top7	961.865	6.841
	MNL		6.984

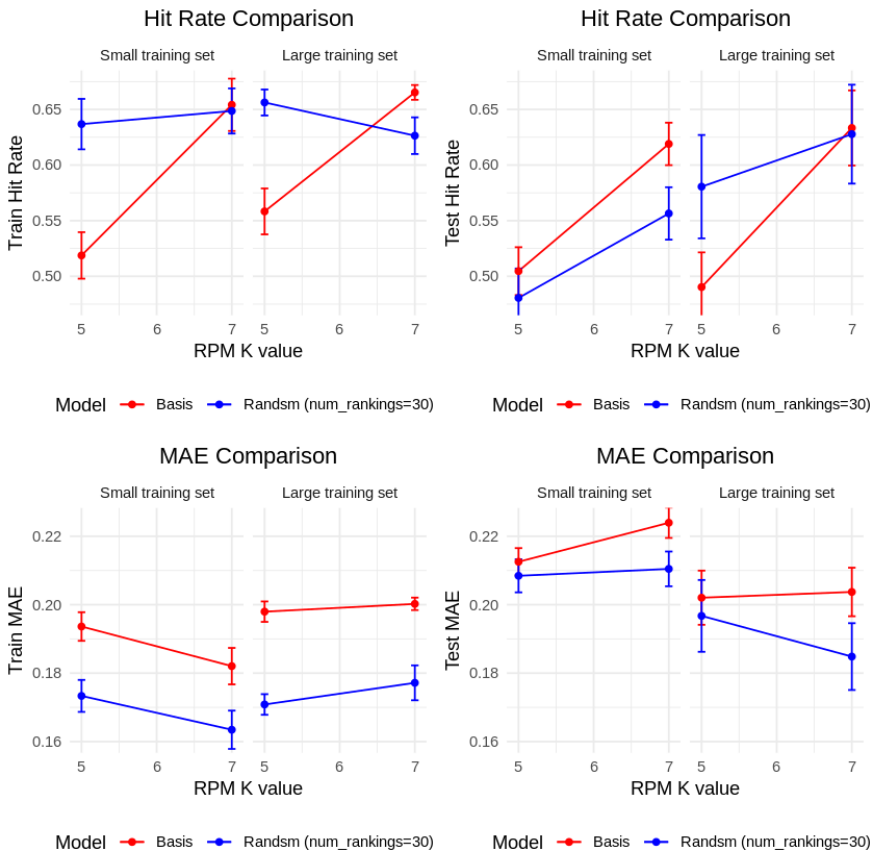

Figure 8 **(Color online) The average hit rate (top row) and MAE (bottom row) in the training (left column) and test set (right column) for RPM based on random sampling with** 30 **rankings (blue line) and the information matrix column basis method (red line) for two different levels of parameter** k **(ranking depth). In each graph, the results are presented for the small training (left) and large training set (right) conditions. The vertical lines represent bootstrap standard errors with** 1000 **replications.**

7. Simulation Study

To better understand the RPM performance, we conducted a computational experiment using synthetic data generated under various configurations of the setting parameters. For data generation, we employed a random lexicographic model (Tversky 1972, Kohli and Jedidi 2007), a flexible non-compensatory choice model capable of capturing violations of the IIA and various forms of context-dependent preferences, such as the decoy and compromise effects. Moreover, this model closely reflects the decision-making processes observed in many real-world application domains.

To generate the choice data, we first draw attribute weights from a Dirichlet distribution, $(\alpha_1, \ldots, \alpha_M) \sim \text{Dir}(\mathbf{1}_M)$, where $\mathbf{1}_M$ denotes an M-dimensional vector of ones. As variation in the Dirichlet-distributed vector $(\alpha_1, \ldots, \alpha_M)$ varies, one or a few attributes tend to become asymmetrically dominant, thereby varying context effects in the generated choice data following a random

lexicographic rule. For each choice task S, attribute priorities are generated according to the following rule:

$$\text{Priority}_m = \log(\alpha_m) - \log(-\log(q_m)), \tag{15}$$

where $q_m \sim \text{Uniform}(0, 1)$. The resulting ordered priority values determine the sequence in which alternatives are screened for that particular choice task S. A new priority sequence is generated independently for each choice task by drawing a new uniform random number q_m, for each attribute m, and applying Eq. (15).

In each iteration of our simulation study, we randomly generated 25 menus. Each option in the menus is randomly generated based on $M = 4$ or 6 attributes, each with five levels. Menus are constructed to ensure the absence of dominated alternatives in each menu. For each menu, we randomly generated 40 choice tasks according to the random lexicographic model described above. Therefore, we generated a total of 1000 choice data in each replication. We then randomly partitioned the menus into the training and test sets. The training set included only the choice data from menus assigned to the training condition. We considered training sets of the size 15 and 20 out of the 25 menus. For each setting, we estimate different variants of RPM using the top-k method with $k = 5$ or 7, and employing either the orthogonal basis or the random sampling method. In the random sampling method, the RPM support size is always restricted to $|\mathcal{P}| = 50$. For each setting, we repeated the process 50 times.

Our results for hit rate and MAE performances are presented in Table 7 and Figure 9 and 10, respectively. Moreover, Table 8 reports the average percentage improvement in RPM compared to the baseline MNL, and their bootstrap 95% confidence intervals. Our results confirm findings from the empirical analysis in the previous section.

Table 7 Mean Model Performance by Number of Attributes and Training Size.

Num. Att.	Train Size	Model Type	In-sample		Holdout	
			Hit Rate (%)	MAE	Hit Rate (%)	MAE
Few	Small	orthogonal basis and top5	**68.0**	0.157	**60.7**	**0.183**
		orthogonal basis and top7	57.6	0.193	55.3	0.204
		random and top5	**68.0**	**0.152**	57.8	0.185
		random and top7	56.6	0.208	55.8	0.219
		MNL	50.9	0.293	50.3	0.300
	Large	orthogonal basis and top5	**67.8**	**0.146**	61.1	**0.178**
		orthogonal basis and top7	58.9	0.190	58.9	0.205
		random and top5	65.9	0.158	**61.8**	0.186
		random and top7	56.1	0.193	56.7	0.205
		MNL	55.9	0.267	51.9	0.280
Many	Small	orthogonal basis and top5	**69.9**	**0.133**	**60.5**	**0.158**
		orthogonal basis and top7	60.6	0.163	53.2	0.173
		random and top5	67.7	0.146	58.0	0.177
		random and top7	62.8	0.169	58.5	0.186
		MNL	61.8	0.201	55.5	0.222
	Large	orthogonal basis and top5	**70.0**	**0.144**	**69.0**	0.162
		orthogonal basis and top7	62.8	0.155	63.3	**0.159**
		random and top5	65.4	**0.144**	58.3	0.167
		random and top7	62.4	0.161	54.7	0.175
		MNL	60.8	0.202	56.0	0.227

Moreover, we compared the total number of rankings in each RPM type with the number of rankings that obtained a nonzero probability value. Our results, presented in Figure 11, show a substantial difference between these two, suggesting that RPM holds a strongly sparsity property. In other words, RPM tends to produce results that are parsimonious. This built-in regularization property is likely a key factor in the model's good performance in holdout predictions.

8. Conclusions

We introduced the Random Preference Model, a flexible, nonparametric, and tractable model of stochastic choice that unifies inferential uncertainty (from limited data) with behavioral randomness (from cognitive factors or aggregate data). Unlike conventional utility-based models that rely on strong parametric assumptions, the RPM describes choices directly as probability distributions over rankings. This provides the model with the flexibility to accommodate context-dependent choice behaviors, such as the compromise effect and violations of independence from irrelevant alternatives, while preserving core axioms like regularity.

While RUM describes choices using a distribution over the utility space, and rank-based models describe choices based on a distribution over rankings, RPM describes choices based on a probability distribution over rankings where the probability of a given ranking is related to a distribution over

Table 8 Summary of Mean Percentage Improvement (RPM over MNL) and 95% Bootstrap Confidence Intervals with 1000 replications, in the computational experiment.

Training Size	Num Att	RPM Type	Metric	Data Set	Mean Improvement (%)	95% CI Lower	95% CI Upper
15 menus	4	orthogonal basis and top5	Hit Rate	Test	**92.4**	43.5	138.0
15 menus	4	orthogonal basis and top5	Hit Rate	Train	**92.3**	36.8	133.0
15 menus	4	orthogonal basis and top5	MAE	Test	**27.8**	16.3	41.1
15 menus	4	orthogonal basis and top5	MAE	Train	**38.0**	28.2	47.5
15 menus	4	orthogonal basis and top7	Hit Rate	Test	**46.2**	5.5	76.4
15 menus	4	orthogonal basis and top7	Hit Rate	Train	**57.1**	5.2	96.4
15 menus	4	orthogonal basis and top7	MAE	Test	**18.6**	5.3	35.2
15 menus	4	orthogonal basis and top7	MAE	Train	**17.2**	2.0	35.3
15 menus	4	random and top5	Hit Rate	Test	18.3	-3.3	36.9
15 menus	4	random and top5	Hit Rate	Train	**91.3**	18.6	145.7
15 menus	4	random and top5	MAE	Test	**25.7**	18.1	32.8
15 menus	4	random and top5	MAE	Train	**38.6**	30.8	46.7
15 menus	4	random and top7	Hit Rate	Test	**85.2**	18.9	138.5
15 menus	4	random and top7	Hit Rate	Train	**34.3**	8.4	58.0
15 menus	4	random and top7	MAE	Test	**16.7**	4.3	32.3
15 menus	4	random and top7	MAE	Train	**15.0**	2.9	28.7
15 menus	6	orthogonal basis and top5	Hit Rate	Test	**22.3**	1.9	43.0
15 menus	6	orthogonal basis and top5	Hit Rate	Train	**25.5**	3.9	45.4
15 menus	6	orthogonal basis and top5	MAE	Test	**17.8**	9.3	27.0
15 menus	6	orthogonal basis and top5	MAE	Train	**17.5**	5.9	30.7
15 menus	6	orthogonal basis and top7	Hit Rate	Test	17.9	-13.4	42.8
15 menus	6	orthogonal basis and top7	Hit Rate	Train	37.6	-11.4	73.6
15 menus	6	orthogonal basis and top7	MAE	Test	1.0	-15.2	19.3
15 menus	6	orthogonal basis and top7	MAE	Train	1.6	-12.5	17.8
15 menus	6	random and top5	Hit Rate	Test	**48.4**	5.6	81.0
15 menus	6	random and top5	Hit Rate	Train	**35.5**	0.9	63.0
15 menus	6	random and top5	MAE	Test	8.2	-2.1	18.9
15 menus	6	random and top5	MAE	Train	6.3	-10.1	27.6
15 menus	6	random and top7	Hit Rate	Test	**35.9**	9.9	62.2
15 menus	6	random and top7	Hit Rate	Train	31.5	-1.7	60.1
15 menus	6	random and top7	MAE	Test	2.9	-8.6	15.2
15 menus	6	random and top7	MAE	Train	9.6	-1.4	21.1
20 menus	4	orthogonal basis and top5	Hit Rate	Test	**46.3**	15.3	71.7
20 menus	4	orthogonal basis and top5	Hit Rate	Train	**98**	28.6	150.6
20 menus	4	orthogonal basis and top5	MAE	Test	**22.8**	7.9	41.8
20 menus	4	orthogonal basis and top5	MAE	Train	**36.4**	25	48.0
20 menus	4	orthogonal basis and top7	Hit Rate	Test	14.3	-6.9	35.3
20 menus	4	orthogonal basis and top7	Hit Rate	Train	**36.2**	0.7	65.0
20 menus	4	orthogonal basis and top7	MAE	Test	12.5	-2.2	30.8
20 menus	4	orthogonal basis and top7	MAE	Train	**14.7**	3.2	27.1
20 menus	4	random and top5	Hit Rate	Test	**41.1**	11.6	66.7
20 menus	4	random and top5	Hit Rate	Train	**45.6**	7.2	76.1
20 menus	4	random and top5	MAE	Test	**16.3**	3.2	31.6
20 menus	4	random and top5	MAE	Train	**26.4**	14.4	41.2
20 menus	4	random and top7	Hit Rate	Test	15.9	-11.9	39.3
20 menus	4	random and top7	Hit Rate	Train	**22.8**	2.3	42.8
20 menus	4	random and top7	MAE	Test	**17.3**	5.6	29.4
20 menus	4	random and top7	MAE	Train	**15.7**	4.1	28.1
20 menus	6	orthogonal basis and top5	Hit Rate	Test	24.9	-1.0	44.5
20 menus	6	orthogonal basis and top5	Hit Rate	Train	**23.9**	4.7	40.0
20 menus	6	orthogonal basis and top5	MAE	Test	8.5	-3.8	21.1
20 menus	6	orthogonal basis and top5	MAE	Train	6.9	-5.8	19.5
20 menus	6	orthogonal basis and top7	Hit Rate	Test	**33.9**	3.5	58.6
20 menus	6	orthogonal basis and top7	Hit Rate	Train	**32**	7.6	52.7
20 menus	6	orthogonal basis and top7	MAE	Test	**20.8**	7.8	33.1
20 menus	6	orthogonal basis and top7	MAE	Train	**17.6**	6.4	28.5
20 menus	6	random and top5	Hit Rate	Test	**26.3**	4.3	48.9
20 menus	6	random and top5	Hit Rate	Train	**23.3**	1.3	42.1
20 menus	6	random and top5	MAE	Test	4.5	-12.0	23.3
20 menus	6	random and top5	MAE	Train	**16.1**	6.3	27.0
20 menus	6	random and top7	Hit Rate	Test	31.2	-11.3	65.8
20 menus	6	random and top7	Hit Rate	Train	18.4	-6.7	37.3
20 menus	6	random and top7	MAE	Test	2.5	-12.6	18.8
20 menus	6	random and top7	MAE	Train	3.7	-8.4	17.9

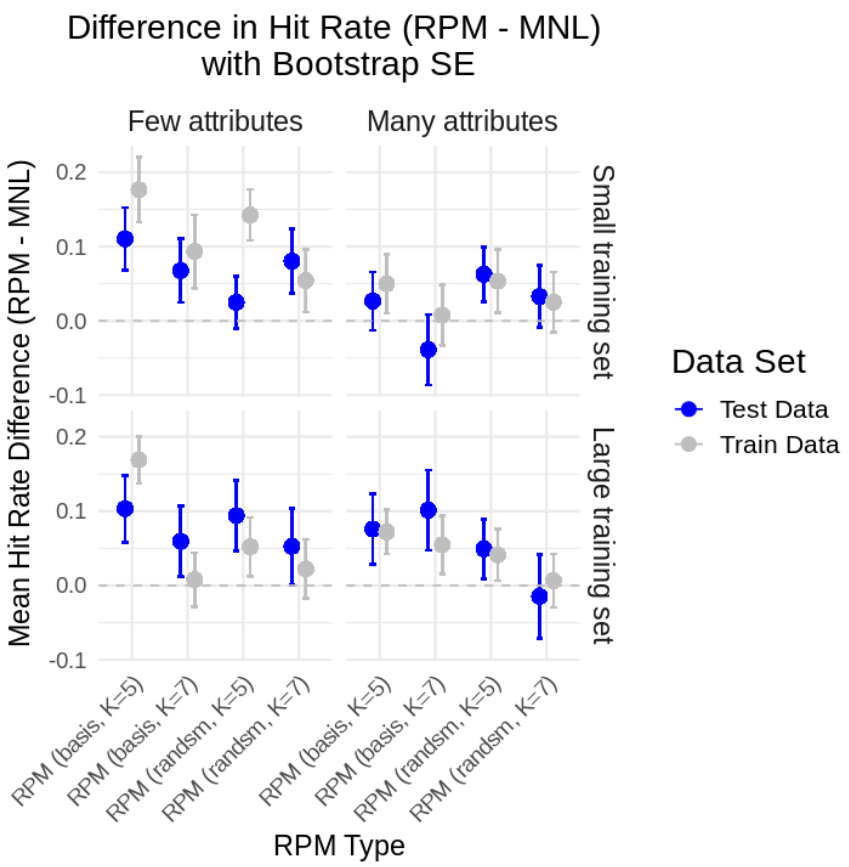

Figure 9 **(Color online) Mean hit rate difference between RPM and MNL in training (gray) and test (blue) sets for small (top) and large (bottom) training sets with few (left) and many (right) attributes, and for different RPM types, in the computational experiment. The vertical lines represent bootstrap standard errors with** 1000 **replications.**

the subspace of utility functions consisting of all utility functions compatible with this ranking. A nested distribution over utility parameters within each ranking further captures intra-ranking variation. Thus, for each ranking, we construct an empirical distribution over the compatible utility subspace such that the expected log likelihood obtained by that distribution aligns with the observed ranking probability. This construction relates the ranking distribution to utility representations, enabling utility-based inference while preserving the rank-based foundation of the model.

To estimate the model, we propose a two-stage procedure: first, estimating the ranking probabilities and global concentration parameter using a constrained optimization problem that ensures the existence of a utility representation; second, identifying within-ranking concentration parameters through integration over parameter subspaces and recovering distributions over utility subspaces. Monte Carlo sampling, initiated from the centroid of each subspace, is used to approximate these integrals. To manage the factorial complexity of the ranking space, we introduced two approxi-

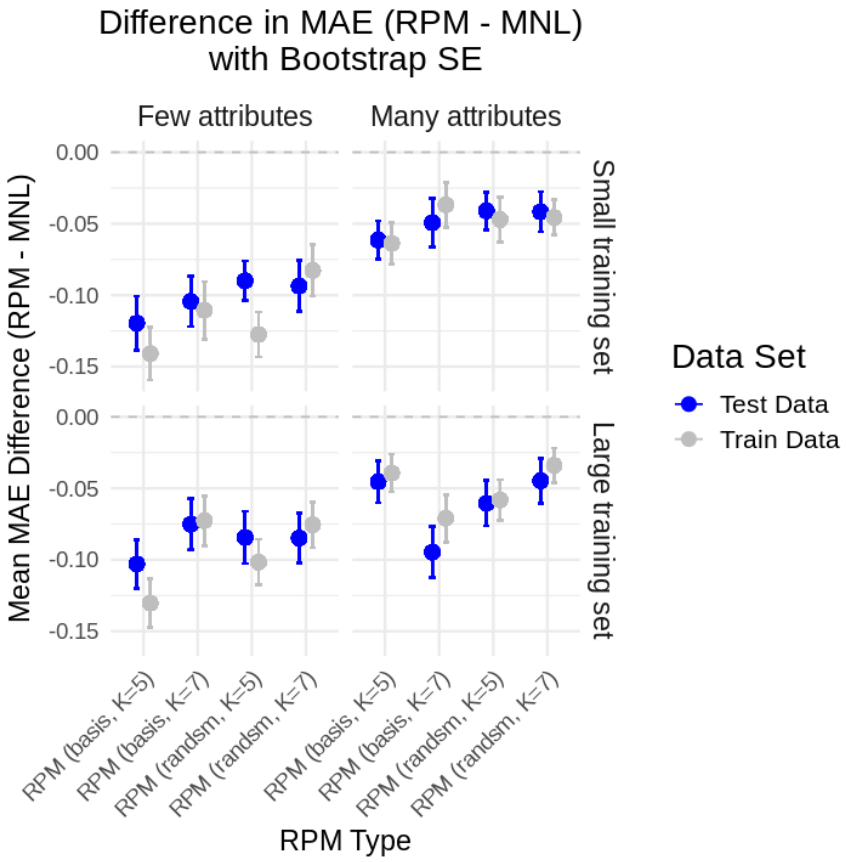

Figure 10 **(Color online) Mean absolute error (MAE) difference between RPM and MNL in training (gray) and test (blue) sets for small (top) and large (bottom) training sets with few (left) and many (right) attributes, and for different RPM types, in the computational experiment. The vertical lines represent bootstrap standard errors with** 1000 **replications.**

mation methods: restricting support to a subset of rankings (via sampling or linear independence) and limiting rankings to top-k alternatives (partial permutations). These strategies enable scalable estimation without substantial loss of performance.

We illustrated the RPM's ability to recover underlying preferences and predict choice behavior under stochastic noise in the prototype study. Even when the observed data included inconsistencies with the true utility-generating process, the full-support RPM demonstrated strong performance in estimating choice probabilities and identifying representative preferences. It successfully reconstructed attribute weights with high accuracy and showed a clear advantage over standard models such as the multinomial logit. Approximate variants of the method – those using restricted support or partial permutations – achieved comparable predictive performance with significantly reduced computational requirements. While these simplifications introduced some loss in precision, the core structure of preferences and key model parameters remained well-recovered.

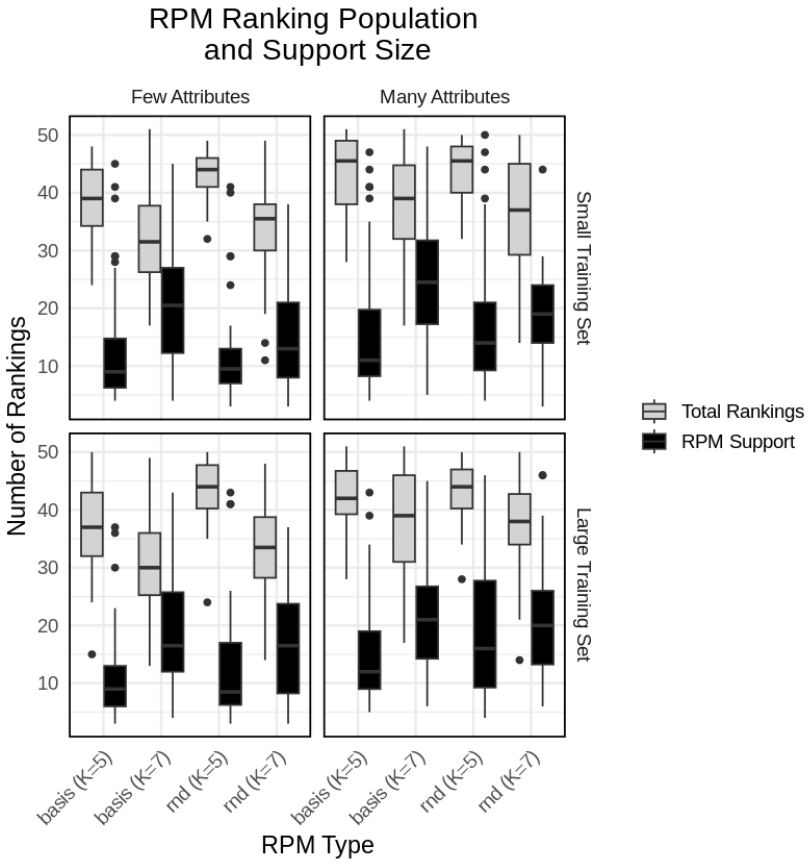

Figure 11 **The average ranking population (grey) and the number of rankings with non-zero probability (black) for small (top) and large (bottom) training sets with few (left) and many (right) attributes, and for different RPM types, in the computational experiment.**

In the empirical study using real-world conjoint data, RPM-based models consistently outperformed the MNL benchmark across various training set sizes. RPM achieved superior predictive accuracy, both in estimating choice probabilities and in identifying the most frequently chosen options. Among the RPM variants, models constructed using a basis of the information matrix offered the best predictive performance, especially when training data was limited. Random sampling methods also performed well, particularly when ranking depth was sufficient and the number of sampled rankings was large. The results showed that increasing the ranking depth improved prediction accuracy across RPM types, with the basis method benefiting more from deeper rankings. In terms of estimation error, the RPM approaches performed similarly as training data increased. While RPM methods required greater computational effort than MNL, the estimation remained tractable and scalable.

While this paper lays the operational foundation for RPM, further research is needed to enhance its theoretical depth, computational scalability, and practical applicability. First, our findings indicate that the partial permutation approximation, which limits rankings to the top-k alternatives, substantially simplifies the RPM without sacrificing predictive accuracy. This result motivates a deeper theoretical investigation into when a stochastic choice function admits a partial permutation representation for a given $k < N$. To date, this remains a largely unexplored area in stochastic choice theory. Future work should also explore how the choice of k interacts with problem characteristics – such as the number of alternatives, attribute dimensionality, or the level of stochastic noise – to influence model performance and generalization. Moreover, treating the ranking depth k as an unknown hyperparameter and estimating it within a *Bayesian framework* is also promising. This would allow the model to adaptively balance expressiveness and complexity based on the observed data.

Second, the computational bottleneck in RPM estimation lies in sampling from the parameter subspaces Θ^P. Scalability may be improved by imposing additional structure on the preference distribution. One favorable direction is the development of a *unimodal RPM*, where the probability of each ranking decays with its distance from an unknown modal ranking. This formulation reduces estimation to identifying the modal ranking and a concentration parameter, offering a more parsimonious representation. The model reflects the behavioral intuition that a decision maker typically relies on a dominant preference order, occasionally deviating toward nearby preference orderings.

Third, when the number of alternatives is large, a key challenge lies in constructing a restricted support set $\mathcal{P}$ that maintains model fidelity while reducing computational complexity. While this paper explores random sampling and column space basis methods, an alternative is to *embed support construction within the estimation process itself*. For example, auxiliary binary variables x_{ijr} can be introduced to encode ranking relationships (e.g., whether item i is ranked above item j in ranking r). These variables, subject to transitivity and consistency constraints, would enable the model to directly construct feasible rankings during estimation. Though this approach introduces a mixed-integer optimization problem, recent advances in discrete optimization and column generation, e.g., van Ryzin and Vulcano 2015, offer promising computational tools for this task.

References

Abdellaoui M (2000) Parameter-free elicitation of utility and probability weighting functions. *Management Science* 46(11):1497–1512.

Aboutaleb YM, Ben-Akiva M, Jaillet P (2020) Learning Structure in Nested Logit Models. URL https://arxiv.org/abs/2008.08048.

Afriat SN (1967) The construction of utility functions from expenditure data. *International Economic Review* 8(1):67–77.

Agranov M, Ortoleva P (2017) Stochastic choice and preferences for randomization. *Journal of Political Economy* 125(1):40–68.

Allen R, Rehbeck J (2023) Revealed stochastic choice with attributes. *Economic Theory* 75(1):91–112.

Alptekinoğlu A, Semple JH (2016) The exponomial choice model: A new alternative for assortment and price optimization. *Operations Research* 64(1):79–93.

Aouad A, Farias V, Levi R (2021) Assortment optimization under consider-then-choose choice models. *Management Science* 67(6):3368–3386.

Aouad A, Feldman J, Segev D (2023) The exponomial choice model for assortment optimization: An alternative to the mnl model? *Management Science* 69(5):2814–2832.

Apesteguia J, Ballester MA, Lu J (2017) Single-crossing random utility models. *Econometrica* 85(2):661–674.

Ariely D (2008) *Predictably Irrational: The Hidden Forces That Shape Our Decisions* (New York: HarperCollins).

Bai Y, Feldman J, Topaloglu H, Wagner L (2024) Assortment optimization under the multinomial logit model with utility-based rank cutoffs. *Operations Research* 72(4):1453–1474.

Baldiga KA, Green JR (2013) Assent-maximizing social choice. *Social Choice and Welfare* 40(2):439–460.

Barberá S, Pattanaik PK (1986) Falmagne and the rationalizability of stochastic choices in terms of random orderings. *Econometrica: Journal of the Econometric Society* 707–715.

Bell DE, Keeney RL, Little JD (1975) A market share theorem. *Journal of Marketing Research* 12(2):136–141.

Ben-Akiva M, Bierlaire M (1999) Discrete choice methods and their applications to short term travel decisions. *Handbook of Transportation Science*, 5–33 (Springer).

Ben-Akiva M, Lerman SR (1985) *Discrete Choice Analysis: Theory and Application to Travel Demand* (Cambridge, MA: MIT Press).

Ben-Akiva ME (1973) *Structure of passenger travel demand models.* Ph.D. thesis, Massachusetts Institute of Technology.

Berbeglia G, Venkataraman A (2023) The generalized stochastic preference choice model. URL https://arxiv.org/abs/1803.04244.

Bertsimas D, Mišic VV (2015) Data-driven assortment optimization. *Working paper, MIT Sloan School* .

Bertsimas D, Mišić VV (2019) Exact first-choice product line optimization. *Operations Research* 67(3):651–670.

Bettman J, Luce M, Payne J (1998) Constructive consumer choice processes. *Journal of Consumer Research* 25(3):187–217.

Blanchet J, Gallego G, Goyal V (2016) A Markov chain approximation to choice modeling. *Operations Research* 64(4):886–905.

Bleichrodt H, Pinto JL (2000) A parameter-free elicitation of the probability weighting function in medical decision analysis. *Management Science* 46(11):1485–1496.

Block HD, Marschak J (1959) Random orderings and stochastic theories of response. Discussion Paper 66, Cowles Foundation for Research in Economics, Yale University, New Haven, CT.

Bordalo P, Gennaioli N, Shleifer A (2013) Salience and consumer choice. *Journal of Political Economy* 121(5):803–843.

Brooks S, Gelman A, Jones G, Meng XL (2011) *Handbook of Markov Chain Monte Carlo* (CRC press).

Chambers CP, Echenique F (2016) *Revealed Preference Theory*, volume 56 of *Econometric Society Monographs* (Cambridge, UK: Cambridge University Press).

Chambers CP, Echenique F, Lambert NS (2018) Preference identification. URL `https://arxiv.org/abs/1807.11585`.

de Bekker-Grob EW, Ryan M, Gerard K (2012) Discrete choice experiments in health economics: a review of the literature. *Health Economics* 21(2):145–172.

Désir A, Goyal V, Jiang B, Xie T, Zhang J (2024) Robust assortment optimization under the Markov chain choice model. *Operations Research* 72(4):1595–1614.

Falmagne JC (1978) A representation theorem for finite random scale systems. *Journal of Mathematical Psychology* 18(1):52–72.

Farias V, Jagabathula S, Shah D (2020) Inferring Sparse Preference Lists from Partial Information. *Stochastic Systems* 10(4):335–360.

Farias VF, Jagabathula S, Shah D (2013) A nonparametric approach to modeling choice with limited data. *Management Science* 59(2):305–322.

Feldman JB, Topaloglu H (2017) Revenue management under the Markov chain choice model. *Operations Research* 65(5):1322–1342.

Feng G, Li X, Wang Z (2017) On the relation between several discrete choice models. *Operations Research* 65(6):1516–1525.

Fishburn PC (1970) *Utility Theory for Decision Making*. Publications in Operations Research, Vol. 18 (New York: John Wiley & Sons).

Fishburn PC (1992) Induced binary probabilities and the linear ordering polytope: A status report. *Mathematical Social Sciences* 23(1):67–80.

Fishburn PC (1998) Stochastic utility. Barberá S, Hammond PJ, Seidl C, eds., *Handbook of Utility Theory: Volume 1, Principles*, 273–318 (Dordrecht: Kluwer Academic Publishers).

Fishburn PC, Falmagne JC (1989) Binary choice probabilities and rankings. *Economics Letters* 31(2):113–117.

Frederick S, Lee L, Baskin E (2014) The limits of attraction. *Journal of Marketing Research* 51(4):487–507.

Gallego G, Ratliff R, Shebalov S (2015) A general attraction model and sales-based linear program for network revenue management under customer choice. *Operations Research* 63(1):212–232.

Ghaderi M (2017) *Preference Disaggregation: Towards an Integrated Framework*. Ph.D. thesis, ESADE Business School.

Ghaderi M, Kadziński M (2021) Incorporating uncovered structural patterns in value functions construction. *Omega* 99:102203.

Ghaderi M, Ruiz F, Agell N (2017) A linear programming approach for learning non-monotonic additive value functions in multiple criteria decision aiding. *European Journal of Operational Research* 259(3):1073–1084.

Hausman DM (2011) *Preference, Value, Choice, and Welfare* (New York: Cambridge University Press).

Hensher DA (1994) Stated preference analysis of travel choices: the state of practice. *Transportation* 21:107–133.

Huber J, Payne JW, Puto C (1982) Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research* 9(1):90–98.

Jacquet-Lagreze E, Siskos J (1982) Assessing a set of additive utility functions for multicriteria decision-making, the uta method. *European journal of operational research* 10(2):151–164.

Jagabathula S (2014) Assortment optimization under general choice. *Available at SSRN 2512831* .

Kadziński M, Ghaderi M, Dabrowski M (2020) Contingent preference disaggregation model for multiple criteria sorting problem. *European Journal of Operational Research* 281(2):369–387.

Kadziński M, Greco S, Słowiński R (2012) Extreme ranking analysis in robust ordinal regression. *Omega* 40(4):488–501.

Keeney RL, Raiffa H (1993) *Decisions with multiple objectives: preferences and value trade-offs* (Cambridge university press).

Kivetz R, Netzer O, Srinivasan V (2004) Alternative models for capturing the compromise effect. *Journal of Marketing Research* 41(3):237–257.

Kohli R, Jedidi K (2007) Representation and inference of lexicographic preference models and their variants. *Marketing Science* 26(3):380–399.

Kohli R, Jedidi K (2017) Relation between eba and nested logit models. *Operations Research* 65(3):621–634.

Lichtenstein S, Slovic P, eds. (2006) *The Construction of Preference* (New York, NY: Cambridge University Press).

Louviere JJ, Hensher DA, Swait JD (2000) *Stated Choice Methods: Analysis and Applications* (Cambridge, UK: Cambridge University Press).

Louviere JJ, Pihlens D, Carson R (2011) Design of discrete choice experiments: a discussion of issues that matter in future applied research. *Journal of Choice Modelling* 4(1):1–8.

Luce RD (1959) *Individual Choice Behavior: A Theoretical Analysis* (New York: Wiley).

McFadden D (1973) *Conditional logit analysis of qualitative choice behavior*, 105–142 (Academic Press).

McFadden D (1974) The measurement of urban travel demand. *Journal of Public Economics* 3(4):303–328.

McFadden D (1977) Modelling the choice of residential location. Discussion Paper 477, Cowles Foundation for Research in Economics, Yale University, New Haven, CT.

McFadden D (2001) Economic choices. *American Economic Review* 91(3):351–378.

McFadden D, Richter M (1970) Revealed stochastic preference. *Unpublished manuscript, Department of Economics, University of California, Berkeley* .

McFadden D, Richter MK (1990) Stochastic rationality and revealed stochastic preference. *Preferences, Uncertainty, and Optimality, Essays in Honor of Leo Hurwicz, Westview Press: Boulder, CO* 161–186.

Miller KM, Hofstetter R, Krohmer H, Zhang ZJ (2011) How should consumers' willingness to pay be measured? An empirical comparison of state-of-the-art approaches. *Journal of Marketing Research* 48(1):172–184.

Richter MK (1966) Revealed preference theory. *Econometrica: Journal of the Econometric Society* 635–645.

Rooderkerk RP, Van Heerde HJ, Bijmolt TH (2011) Incorporating context effects into a choice model. *Journal of Marketing Research* 48(4):767–780.

Rubinstein A (2012) *Lecture Notes in Microeconomic Theory: The Economic Agent* (Princeton, NJ: Princeton University Press), 2nd edition.

Salant Y, Rubinstein A (2008) (a, f): choice with frames. *The Review of Economic Studies* 75(4):1287–1296.

Samuelson PA (1938) A note on the pure theory of consumer's behaviour. *Economica* 5(17):61–71.

Samuelson W, Zeckhauser R (1988) Status quo bias in decision making. *Journal of Risk and Uncertainty* 1(1):7–59.

Simonson I (2014) Vices and virtues of misguided replications: The case of asymmetric dominance. *Journal of Marketing Research* 51(4):514–519.

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

Srinivasan V, Shocker AD (1973) Linear programming techniques for multidimensional analysis of preferences. *Psychometrika* 38(3):337–369.

Strzalecki T (2017) Stochastic choice. *Hotelling Lectures in Economic Theory, Econometric Society European Meeting* (Econometric Society).

Sturt B (2025) The Value of Robust Assortment Optimization Under Ranking-Based Choice Models. *Management Science* 71(5):4246–4265.

Toubia O, Hauser J, Garcia R (2007) Probabilistic polyhedral methods for adaptive choice-based conjoint analysis: Theory and application. *Marketing Science* 26(5):596–610.

Toubia O, Simester DI, Hauser JR, Dahan E (2003) Fast polyhedral adaptive conjoint estimation. *Marketing Science* 22(3):273–303.

Train KE (2009) *Discrete Choice Methods with Simulation* (New York, NY: Cambridge University Press).

Turansick C (2022) Identification in the random utility model. *Journal of Economic Theory* 203:105489.

Tversky A (1972) Elimination by aspects: A theory of choice. *Psychological Review* 79(4):281–299.

Tversky A, Sattath S (1979) Preference trees. *Psychological Review* 86(6):542–573.

Tversky A, Simonson I (1993) Context-dependent preferences. *Management Science* 39:1179–1189.

van Ryzin G, Vulcano G (2015) A market discovery algorithm to estimate a general class of nonparametric choice models. *Management Science* 61(2):281–300.

van Ryzin G, Vulcano G (2017) An expectation-maximization method to estimate a rank-based choice model of demand. *Operations Research* 65(2):396–407.

Wu C, Cosguner K (2020) Profiting from the decoy effect: A case study of an online diamond retailer. *Marketing Science* 39(5):974–995.