



# Rank-heterogeneous Preference Models for School Choice

Amel Awadelkarim  
ameloa@stanford.edu  
Stanford University  
Stanford, CA, USA

Arjun Seshadri  
sesarjun@amazon.com  
Amazon  
San Francisco, CA, USA

Itai Ashlagi  
iashlagi@stanford.edu  
Stanford University  
Stanford, CA, USA

Irene Lo  
ilo@stanford.edu  
Stanford University  
Stanford, CA, USA

Johan Ugander  
jugander@stanford.edu  
Stanford University  
Stanford, CA, USA

## ABSTRACT

School choice mechanism designers use discrete choice models to understand and predict families' preferences. The most widely-used choice model, the multinomial logit (MNL), is linear in school and/or household attributes. While the model is simple and interpretable, it assumes the ranked preference lists arise from a choice process that is uniform throughout the ranking, from top to bottom. In this work, we introduce two strategies for rank-heterogeneous choice modeling tailored for school choice. First, we adapt a context-dependent random utility model (CDM), considering down-rank choices as occurring in the context of earlier up-rank choices. Second, we consider stratifying the choice modeling by rank, regularizing rank-adjacent models towards one another when appropriate. Using data on household preferences from the San Francisco Unified School District (SFUSD) across multiple years, we show that the contextual models considerably improve our out-of-sample evaluation metrics across all rank positions over the non-contextual models in the literature. Meanwhile, stratifying the model by rank can yield more accurate first-choice predictions while down-rank predictions are relatively unimproved. These models provide performance upgrades that school choice researchers can adopt to improve predictions and counterfactual analyses.

## CCS CONCEPTS

- Information systems → Rank aggregation; • Applied computing → Economics.

## KEYWORDS

school choice, discrete choice, preference modeling, ranking models

### ACM Reference Format:

Amel Awadelkarim, Arjun Seshadri, Itai Ashlagi, Irene Lo, and Johan Ugander. 2023. Rank-heterogeneous Preference Models for School Choice. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23)*, August 6–10, 2023, Long Beach, CA, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3580305.3599484>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

KDD '23, August 6–10, 2023, Long Beach, CA, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 979-8-4007-0103-0/23/08...\$15.00  
<https://doi.org/10.1145/3580305.3599484>

## 1 INTRODUCTION

Large school districts around the world employ *school choice* mechanisms to assign students to K–12 schools. In many of these systems, families submit ranked preference lists over school programs to their district, and the district in turn assigns children to schools via a centralized mechanism. School choice researchers employ *discrete choice* models, statistical models of choices made from slates of discrete options, to describe the preference-generation process by breaking a ranking into a sequence of choices from dwindling choice sets.

Such models are useful for explanation, indirectly identifying the most influential school characteristics in the decision-making of families, saving time and resources in surveying families. They can also be used for forecasting and planning potential changes in the district offerings. Finally, these models are also central to evaluating changes in school choice mechanisms themselves, as policymakers propose changes to assignment mechanisms with the hope of improving district outcomes. In the latter contexts, these models play a role in simulating preferences and assignments, and/or evaluating the resulting welfare of assignment under the proposed mechanism. Put simply, better preference models lead to better school choice analyses, and better analysis lead to better childhood educational outcomes.

The widely-used ranked preference models in this space, including the Plackett–Luce “exploded logit” model [22, 28], model the process of constructing a preference ranking as a series of independent discrete choices (conditional multinomial logit (MNL) in the case of Plackett–Luce) based on school, program, and household attributes. While many such models are simple and interpretable, there is long-standing evidence in the discrete choice literature for ranking behavior that is *rank-heterogeneous*, meaning that the sequence of choices are driven by different considerations as individuals work down a preference list [10, 12, 16]. The criteria an agent uses for selecting top-ranked alternatives may differ from those at lower ranks, either due to true preference shifts or behavioral mechanisms such as decision fatigue.

In this work, we present and evaluate two strategies for incorporating rank-heterogeneity in choice models for school choice. One strategy achieves heterogeneity through a sequential dependence using *context effects*, while the other relies on regularized model stratification.

*Rank-heterogeneity via context effects.* *Context effects* describe the influence of a particular decision context, including the available

or previously-chosen options, on a individual’s relative preferences between alternatives. We adapt a previous model of context effects, the context-dependent random utility model (CDM) [30], to the school ranking setting. The CDM has been used to study ranked preferences [31] by decomposing the ranking process as a series of choices in the context of the dwindling set of items yet to be chosen. We consider a variation of the CDM more natural to the school choice setting: modeling the ranking instead as a series of choices in the context of the *already chosen* items. Surprisingly, we show that the two modeling approaches (respectively, forward-dependence and backwards-dependence) are equivalent, and opt to use the latter variation when interpreting our results.

*Rank-heterogeneity via model stratification.* An alternative approach to inducing rank-heterogeneity is stratifying the modeling problem by rank position. Simply learning a series of independent models for each rank position, however, can split the data too finely and result in poor generalization. To avoid this pitfall, we apply Laplacian regularization [37] to the independent models, with carefully tuned regularization graphs that bring models of adjacent choices close together.

Incorporating context effects and model stratification are not mutually exclusive, and we also evaluate the combination of both approaches in our analysis. Moreover, we perform a series of ablation studies to demonstrate the independent contributions of each approach. We evaluate these new tools by modeling the preferences for the San Francisco Unified School District (SFUSD) kindergarten programs during the 2017-18 and 2018-19 assignment years. We find that the first strategy (context effects) dramatically lowers out-of-sample negative log likelihood, particularly on down-rank choices, when compared to rank-homogeneous models. The second strategy (model stratification) delivers more accurate prediction in top choices than a rank-homogeneous model—essentially, by modelling them separately—but otherwise does not appear to produce any significant improvements over the non-stratified baseline. Furthermore, we evaluate the performance of our context effect model against a nested MNL model and demonstrate sizable advantages in the school choice setting.

*Outline.* Section 2 introduces notation and definitions used throughout the work. Section 3 explains the SFUSD assignment system, its inputs and outputs, and summarizes the data we use for training and evaluation. In Section 4, we describe the choice models studied in this work, presenting the backwards-dependent context-dependent model (CDM) and the stratified approach with Laplacian regularization. Section 5 addresses identifiability of the models and details our model optimization framework. In Section 6, we present and discuss the performance of our models; Section 7 concludes. Additional discussion and analyses are found in the appendix of our full text.<sup>1</sup>

## 1.1 Related Work

The present work closely relates to various prior works that develop or apply preference models in school choice. Laverde [20] uses an MNL choice model to simulate counterfactual assignments in Boston in 2010–2013, quantifying the role of distance and unequal

access on stated preferences. Agarwal & Somaini [2] develop a procedure for estimating an MNL model in the presence of strategic reporting. Abdulkadiroğlu et al. [1] use MNL models to find links between preferences, school effectiveness and peer quality in New York City in 2003–2013. For an in-depth review of prior applications of preference models in school choice, see Agarwal & Somaini [3].

Meanwhile, many works have studied the relative suitability of different choice models in school choice, evaluating accuracy and prediction errors of preference models. For example, Pathak & Shi [27] examine out-of-sample estimates for three models after a large-scale policy change in Boston. They develop several model evaluation metrics, and we adapt one to our work. Calsamiglia et al. [9] similarly estimate a full choice system and evaluate it out-of-sample using administrative data from 2006 and 2007 school years in Barcelona.

Several prior efforts aim to understand preference heterogeneity between various demographic groups. For example, Laverde [20] estimates MNL models for White, Black, and Hispanic families by including indicator variables for these features in the chosen MNL utility. Hastings et al. [15] apply mixed-logit models [24] to data from Charlotte-Mecklenburg, North Carolina, learning separate model coefficients by race and SES status. In contrast to these examples of heterogeneity *between* groups, the present work focuses instead on preference heterogeneity *within* participants as they assemble their rankings.

Our idea is inspired by prior works in psychology, economics and marketing research, all of which cite inconsistent agent behavior in the assembly of rankings. Under the observation that individuals are generally more careful in reporting their top choices than lower ranked ones, Hausman & Ruud [16] model structured rank-heterogeneity through a common choice model with increasing variance as choosers proceed down the ranks, Chapman & Staelin [10] drop ranked alternatives after a threshold, and Allison & Christakis [4] interact model covariates with indicators for early (top-4) or late (5+) rank choices. Our work extends this last idea by fully stratifying models by rank position of choice, interacting all model parameters with indicators for the first  $k$  ranks. More on our stratification (and regularization) framework in Section 4.

Finally, our work applies recent advances from the discrete choice and preference learning literatures to the school choice domain. The MNL model satisfies the axiom of independence of irrelevant alternatives (IIA), that the relative probability of selecting any item  $j$  over another item  $k$  from choice set  $S$  is independent of the other items in  $S$ . However, this axiom is highly restrictive and often not representative of the true choice process [38, 39]. We adopt strategies for going beyond the independence of irrelevant alternatives (IIA) assumption from Seshadri et al. [30], in turn adapted from Batsell & Polking [7], extending that framework from a previously-studied forward-dependent model of ranking [31] to backwards-dependent ranking. Other recent work extending the CDM include studies of salient features [8] and feature-based context effects [34]; we leave the evaluation of such model extensions as future work. Further, we benchmark the performance of our approaches against the nested MNL model [23], which also goes beyond the restrictive IIA assumption, in Section 6.1.

<sup>1</sup>Full paper with appendices can be found at <https://arxiv.org/abs/2306.01801>.

## 2 CHOICE PRELIMINARIES

We begin by introducing our notation for viewing school choice through the lens of discrete choice. For a specific school year, let  $\mathcal{U} := [m] = \{1, \dots, m\}$  denote the universe set of all offerings, or *alternatives*, in the district, labeled 1 through  $m$ , and let  $n$  be the number of students seeking assignment in the choice system. Throughout this work, we use “household” and “student” interchangeably to represent the decision-maker, as enrollment pertains to the student but rankings are often submitted by caretakers. Further, let  $\text{PO}(\mathcal{U})$  denote the set of all partial orders on the alternatives in  $\mathcal{U}$ . A preference list  $R_i \in \text{PO}(\mathcal{U})$  is household  $i$ ’s partial ranking of the alternatives in  $\mathcal{U}$ , and we denote by  $k_i \leq m$  the length of that ranking. The vector of observable covariates on student  $i$  and offering  $j \in \mathcal{U}$  are given by  $x_{ij}$ , containing demographic, socioeconomic, geographic, and performance-related information on the pair. Then, a *school choice dataset*,  $(\mathcal{D}, X)$ , is defined as the collection of all participating-household’s partial rankings submitted to the district,  $\mathcal{D} = \{R_1, \dots, R_n\}$ , and observed student-program covariates,  $X \in \mathbb{R}^{n \times m \times d}$ , where  $x_{ij}$  is a length- $d$  vector of attributes pertaining to student  $i$  and alternative  $j$ .

To learn a model of rank data, researchers typically transform rankings to choices and then apply discrete choice models such as the MNL, resulting in what is known (equivalently) as the *rank-ordered logit* [16], *exploded-logit* [10, 29], or Plackett-Luce [22, 28] model for rankings, which we present in Section 4. The generality of converting rankings to choice is non-obvious, but the most powerful and widespread transformation is motivated by the theory of L-decomposable ranking distributions [11, 21] (L as in Left). A ranking distribution is said to be L-decomposable if the probability  $P(R)$  of observing ranking  $R = (r_1, \dots, r_k)$  can be decomposed into probabilities of choices from dwindling choice sets, from most to least preferred:

$$P(R) = P(r_1 | \{r_1, \dots, r_k\}) \cdot P(r_2 | \{r_2, \dots, r_k\}) \cdot \dots \cdot P(r_{k-1} | \{r_{k-1}, r_k\}).$$

This unraveling-from-the-left decomposition is sometimes also referred to as *repeated selection* [31]. In the present work, we apply repeated selection to ranking data throughout, simplifying the name of the ranking model to just the enlisted choice model employed after unraveling.

Encoding the unraveled choices as (agent, choice, choice set) triples, the rank data  $\mathcal{D}$  then becomes a choice dataset,  $D$ :

$$D = \bigcup_{R_i \in \mathcal{D}} \bigcup_{j \in [k_i]} (i, r_{ij}, S_{ij}) \quad (1)$$

where  $r_{ij}$  is represents the  $j$ -th selection by agent  $i$  on ranking  $R_i$ , and  $S_{ij} \subseteq \mathcal{U}$  is the slate of available alternatives, or choice set, when choosing position  $j$  of ranking  $R_i$ . The size of the resulting dataset is  $|D| = \sum_{i \in [n]} k_i$ .

To concretely illustrate the decomposition at the level of a data point, given a universe of alternatives  $\mathcal{U} = \{a, b, c, d\}$ , consider a dataset made up of one ranking, by agent 1,  $\mathcal{D} = \{R_1\}$ , where  $R_1 = (b, d, c)$ . Following Eq. (1), the choice dataset becomes

$$\begin{aligned} D &= \{(1, r_{11}, S_{11}), (1, r_{12}, S_{12}), (1, r_{13}, S_{13})\} \\ &= \{(1, b, \{a, b, c, d\}), (1, d, \{a, c, d\}), (1, c, \{a, c\})\}. \end{aligned}$$

## 3 SAN FRANCISCO SCHOOL CHOICE

In this section, we present the assignment process implemented within the San Francisco Unified School District (SFUSD) from 2014 to the present. SFUSD is made up of 130 schools with 150+ unique program offerings. Students enroll in these programs via an annual assignment lottery where families submit ranked preferences over available offerings to the district, and the district tries to honor family choices while satisfying capacity constraints. The algorithm performing this constrained assignment is the student-proposing deferred acceptance algorithm [13]. Participation may occur across all grade levels, but kindergarten enrollment is by far the largest participating group each year, making up over a third of all annual participants. As such, and following suit with many other studies of school choice, we focus solely on kindergarten assignment.

In the face of overly-demanded program seats, the district uses the following priority hierarchy to make assignments:

- (1) **Sibling**: Highest priority. Given to younger siblings of students enrolled at the school.
- (2) **PreK/TK**: Given to students who (1) live in the attendance area of the school (if applicable), and (2) are enrolled in a PreK or TK program at the school itself or in the attendance area of the school (if applicable).
- (3) **Test score area (“CTIP1”)**: Given to students living in neighborhoods with low average test scores. Grants priority across the district, not just to one program or school.
- (4) **Attendance area (AA)**: Given to students living in the attendance area of the school.
- (5) **No priority**: The absence of any of the above priorities.

For each program, a student is considered in the highest priority category for which they qualify. Within each priority tier, ties are broken by next highest tier if applicable, or by a random number,  $v_{ij}$ , drawn uniformly at random for each student-school pair<sup>2</sup>.

Once all submitted preference lists have been exhausted by the matching algorithm, there may be students left without any assignment, for which the district administratively assigns these students to a program not on their list. In this work, as we are solely interested in modeling the preferences submitted by families in the first stage, such assignments fall outside the scope of our analysis.

### 3.1 Dataset

To understand families stated preferences, we study data from both the 2017–18 and 2018–19 school years, principally training models on the 2017–18 data and evaluating out-of-sample on the 2018–19 data. We opted for this train–test split, following other work in school choice [9, 27], to prevent data leakage. This split also mimics real use cases of school choice preference models, where models are used to simulate future years’ outcomes.

Within each year, we collapse all three rounds of stated preference elicitation (instead of focusing only on the first and largest round), to better capture the full district demographics. We exclude programs that are newly offered the 18–19 school year, dropping 804 households from the test dataset, as our model’s fixed effects do

<sup>2</sup>This lottery design is known as the *multiple tie-breaking rule* (MTB) as students receive multiple values, one for each ranked school. By contrast, the single tie-breaking rule (STB) assigns a single lottery value to each student, used across desired schools. For more on the analysis of tie-breaking rules, see [5, 6, 14? ].

**Table 1: Summary statistics of SFUSD dataset, by school year.**

	School year	
	2017-18	2018-19
No. participating households, $n$	5,115	4,329
Total offerings, $m$	154	148
No. unique schools, $n_s$	72	72
No. unique program types, $n_p$	22	19
Avg. length of ranking, $\bar{k}_i$	9.95	7.05
Size of choice dataset, $\sum_i k_i$	49,882	29,810
Percent students CTIP1	16.7%	18.7%

not extrapolate to never-before-seen alternatives. Handling these out-of-distribution preferences is a known limitation of the MNL class of models and an area of future work. Summary statistics for each school year's data are found in Table 1.

The student-program covariates  $X$  used in this work were selected by domain experts at SFUSD:

- Distance: scalar, in miles,
- Square-root distance: scalar, in sqrt. miles,
- Square-root distance  $\times$  CTIP1: scalar, in sqrt. miles,
- Within 0.5 miles: indicator,
- Bus route: indicator for whether the district has bus routes between student ZIP code and school ZIP code,
- Sibling match: indicator for whether the student has one or more sibling(s) already enrolled at the school (not necessarily same program),
- Language match: indicator for whether a language program is in a student's (non-English) home language,
- Attendance area school: indicator for whether the student lives in the attendance area of the school,
- PreK/TK continuation: indicator for whether or not student is enrolled in an SFUSD Pre-K or transitional kindergarten in the same attendance area as or within the school.

We consider the following school-specific features as well, modeled as interacting with the CTIP1-status of the student.

- Average color: state-defined metric quantifying school's absolute performance and improvement in English/language arts, math, chronic absenteeism, and suspension rates. Ordinal color code in each category, encoded as 1-5, and averaged (higher is better),
- Fraction reduced lunch: fraction of the school's population that qualifies for free or reduced-price lunch by the district,
- Before/after school programs: indicator for whether or not school offers before- or after-school programs.

We acknowledge that additional attention to feature engineering can likely yield measurable performance improvements, but we consider the above features adequate and realistic for our purposes, namely evaluating the value of modeling rank-heterogeneity.

## 4 CHOICE MODELS FOR SCHOOL CHOICE

A choice model models probability distributions over subsets of a collection. More formally, let  $\mathcal{S} = \{S : S \subseteq \mathcal{U}, |S| \geq 2\}$  denote the set of all subsets of size at least two of a collection  $\mathcal{U}$ . Let  $P(j|i, S)$  describe, for each agent  $i \in [n]$  and each set  $S \in \mathcal{S}$ , the probability

of agent  $i$  selecting item  $j$  from set  $S$ . Recall from Section 2 that in the SFUSD school choice mechanism,  $n$  households submit partial rankings  $R_1, R_2, \dots, R_n$  with student-program covariates  $X$ . Each partial ranking  $R_i$  is decomposed into choices per Equation (1), obtaining a dataset  $D$  of choices.

We begin by considering a random utility model (RUM) of choice. The utility to agent  $i$  of alternative  $j$  in choice set  $S$  is given by

$$U(j|i, S) = V(j|i, S) + \epsilon_{ij},$$

decomposed into a part labeled  $V$  that is known by the researcher up to some parameters, the *representative utility*, and an unknown part  $\epsilon$  that is treated as random [35]. Under the assumption of independent Gumbel noise  $\epsilon$ , agent  $i$ 's probability of choosing alternative  $j$  from choice set  $S$  is given in closed form by

$$P(j|i, S) = \frac{e^{V(j|i, S)}}{\sum_{k \in S} e^{V(k|i, S)}}, \quad (2)$$

deriving the most ubiquitous RUM—especially in school choice—the conditional multinomial logit (MNL) [21].

Taking the noise instead to be *jointly* Gumbel distributed with correlation yields variations on a mixed MNL [24] or nested MNL [35] model, the latter featuring correlations across pre-specified clusters of alternatives. We benchmark our performance against a nested MNL model in Section 6. Mixed MNL models have performed comparably to ordinary MNL in several prior school choice studies [15, 26], so we do not benchmark against it in this work.

Under the MNL model, the task of the researcher is to define a representative utility function, typically a parametric model, denoted  $V_\theta$ . We select our model from the chosen model class using regularized maximum likelihood, selecting parameters  $\theta$  to minimize

$$\ell(D; \theta) = \ell(D; \theta) + r(\theta), \quad (3)$$

where  $\ell(D; \theta)$  is the negative log-likelihood (NLL) loss,

$$\ell(D; \theta) = -\frac{1}{|D|} \sum_{(i, r_{ij}, S_{ij}) \in D} \log(P_\theta(r_{ij}|i, S_{ij})),$$

$r(\theta) = \lambda \|\theta\|_2$  is the  $\ell_2$  penalty, and  $\lambda$  is the regularization gain.

### 4.1 Basic utilities

At this point, our task is to define the representative utility,  $V$ . Assigning inherent utilities to each alternative,

$$V_\theta(j|i, S) = \delta_j, \quad (4)$$

reduces the model to the basic Plackett-Luce model [21]. Here,  $\delta \in \mathbb{R}^{\tilde{m}}$  where  $\tilde{m} < m$  is defined as the number of unique schools plus the number of unique program-types offered in the district<sup>3</sup>,  $\tilde{m} = n_s + n_p$ . See Table 1 for district summary statistics. We will refer to this model as the **fixed-effect MNL**.

Adding user-alternative specific covariates yields a **linear MNL**,

$$V_\theta(j|i, S) = \delta_j + \beta^T x_{ij}, \quad (5)$$

<sup>3</sup>Each alternative in the choice universe  $j \in \mathcal{U}$  has an associated school,  $s(j)$ , and program type,  $p(j)$ . Example program types are general education, special education, and language program offerings. As such, our fixed effect  $\delta_j$  is actually shorthand for  $\delta_{s(j)} + \delta_{p(j)}$ , reducing degrees of freedom while allowing our models to better generalize to new offerings at existing schools. We refer the interested reader to our code for the exact implementation.

the most common utility structure in the school choice literature. Note that covariates contained in  $x_{ij}$ , detailed in the previous section, are indexed both by student and alternative; school- or program-specific features (only indexed by  $j$ ) are absorbed into the fixed effects  $\delta_j$ , and student-specific features (only indexed by  $i$ ) cancel out in the expression of the MNL choice probability, Eq. (2).

As an auxiliary benchmark, we implement a **nested MNL** model, using the same expression of representative utility as the linear MNL in Eq. (5), for benchmarking in Section 6. In the nested MNL, alternatives are explicitly assigned to one of  $K$  nests (non-overlapping subsets of  $\mathcal{U}$ ), and choice probabilities are defined to be correlated within nests. See online Appendix C.3 for full discussion and presentation of the nested MNL choice probabilities. The nests we implement in this work are ‘Chinese Language’, ‘Filipino Language’, ‘General Education’, ‘Japanese Language’, ‘Korean Language’, ‘Spanish Language’, and ‘Special Education’ offerings. Each nest is associated with a number of unique program offerings —see Table 1 for more details.

The fixed-effect and linear MNL models presented in this section satisfy IIA (see Section 1.1). As a result, they are *rank-homogeneous*, relying on a constant representative utility  $V(j|i, S; \theta)$  for each alternative regardless of when the choice is being made in the ranking process. The contextual choice model that follows does not satisfy IIA and thus leads to rank-heterogeneous choice distributions.

## 4.2 Context effects

The **context-dependent model** (CDM) [31] is our first strategy for incorporating rank-heterogeneity into the traditional MNL models above. The CDM relaxes the strict IIA assumption, initially crafted to model “choice set effects” [36] whereby the slate of alternatives under consideration impacts the choice probabilities of the agent. We modify this modeling framework to suit the school choice sequential ranking problem. Specifically, under the standard CDM, each choice  $j$  from choice set  $S$  occurs within the context of the choice set  $S$ . For our purposes, we generalize this framework to consider the choices as occurring within the context of a generic and possibly different *context set* of alternatives,  $A$ .

The representative utility of this generalized CDM models context effects as a linear dependence between items, interpretable as “push” and “pull” factors, with items in  $A$  pushing and pulling on each alternative in the choice set  $S$ ,

$$V_\theta(j|i, A, S) = \delta_j + \beta^T x_{ij} + \frac{1}{|A|} \sum_{k \in A} u_{jk}, \forall j \in S.$$

When  $A = S \setminus j$  we recover the standard CDM. The parameters  $u_{ij}$  are defined for all  $i, j \in \mathcal{U}$  where  $i \neq j$ . The generalized CDM has the same parameter complexity as standard CDM, requiring  $m(m-1)$  parameters beyond the linear model, arranged in a matrix-like structure  $U \in \mathbb{R}^{m \times m}$ , with undefined diagonal. To reduce the parametric complexity of the model,  $U$  can be factorized as the product of two low-rank matrices,  $U = TC^T$  with  $T, C \in \mathbb{R}^{m \times r}$  serving as *target* and *context* embeddings, respectively, analogous to word2vec-type methods [25]. The low-rank CDM representative utility is then written as

$$V_\theta(j|i, A, S) = \delta_j + \beta^T x_{ij} + \frac{1}{|A|} \sum_{k \in A} t_j^T c_k. \quad (6)$$

We proceed with the factorized form of the CDM in this work. The low-rank CDM introduces a hyperparameter, in the form of the embedding dimension  $r$ ; see Section 5 for a discussion of hyperparameter tuning.

To accompany this change in model, the structure of the data described in Eq. (1) must be generalized to include a generic context set for each choice, resulting in the following choice dataset

$$D = \bigcup_{i \in \mathcal{D}} \bigcup_{j \in [k_i]} (i, r_{ij}, A_{ij}, S_{ij}), \quad (7)$$

where  $A_{ij}$  is the context set when agent  $i$  chose item  $j$ . Table 2 summarizes the representative utilities of the three models—the fixed-effect MNL, linear MNL, and CDM—and their parameters.

*Forward vs. backward-dependence.* In the original formulation of the CDM for rankings [31], the context set was assumed to be the choice set itself,  $A = S \setminus j$ , a formulation we refer to as the *forward-dependent* contextual ranking model.

Considering the generalized CDM above, we consider instead a model where the context set  $A$  is the set of *already-chosen* alternatives. Equivalently, let  $A = \mathcal{U} \setminus S$ , the complement of the current choice set. We introduce this model as the *backward-dependent* contextual ranking model. Rather than modeling context effects between alternatives in the choice set, it measures how well each alternative fits with the choices already made. This conceptual shift is better suited to the psychology of the school choice selection process than the former framing, and yields a more interpretable model in ranking settings where choice sets are large, such as in school choice.

Considering these two different approaches to modeling rankings as a sequence of contextual choices, it seems as though these formulations constitute different model classes. However, we find that the forwards- and backwards-dependent CDM ranking model classes are in fact equivalent, and provide a bijection between the spaces of parameters for both the unfactorized and factorized models. See online Appendix A for proofs of Theorems 1 and 2.

**THEOREM 1.** Let  $\theta_F = \{\delta^F, \beta^F, U^F\}$  denote model parameters of the unfactorized forward-dependent CDM ranking model, and  $\theta_B$  denote those of the backward-dependent model. The forward- and backward-dependent parameters are equivalent under the bijection  $\theta_B = f(\theta_F)$ , where

$$f(\theta) = \left\{ \delta_i + \sum_{j \in \mathcal{U} \setminus i} u_{ij}, \forall i \right\}, \beta, -U$$

The inverse map is the map itself,  $f^{-1} = f$ .

**THEOREM 2.** Let  $\theta_F = \{\delta^F, \beta^F, T^F, C^F\}$  denote model parameters of the low-rank forward-dependent CDM ranking model, and  $\theta_B$  denote those of the low-rank backward-dependent model. These model parameters are equivalent under the bijection  $\theta_B = g(\theta_F)$ , where

$$g(\theta) = \left\{ \delta_i + t_i^T \sum_{j \in \mathcal{U} \setminus i} c_j, \forall i \right\}, \beta, T, -C$$

The inverse map is the map itself,  $g^{-1} = g$ .

In Section 6, as a supplemental analysis, we consider results for truncated top- $k$ -dependent context sets,  $A_{ij} = \{r_{i1}, \dots, r_{il}\}$  for

**Table 2: Summary of models and the number of degrees of freedom,  $N_p$ . Here  $m$  denotes the number of alternatives (school and program pairs) offered by the district,  $\tilde{m} = n_s + n_p$  denotes the total number of unique schools and program types,  $d$  is the length of  $x_{ij}$ , and  $r$  is the embedding dimension of the low-rank CDM.**

Model	$V_\theta(j i, A, S)$	$\theta$	$N_p$
Fixed	$\delta_j$	$\{\delta\}$	$\tilde{m}$
Linear	$\delta_j + \beta^T x_{ij}$	$\{\delta, \beta\}$	$\tilde{m} + d$
CDM	$\delta_j + \beta^T x_{ij} + \frac{1}{ A } \sum_{k \in A} t_j^T c_k$	$\{\delta, \beta, T, C\}$	$\tilde{m} + d + 2rm$

$l = \min(k, j)$ , to evaluate whether a more limited dependence (and thus, simpler model) performs as well as full backward-dependence. We find that even the truncated top-1 CDM—with knowledge only of the agent’s first choice—makes considerable gains over the linear MNL model, but the full (top- $m$ ) backwards-dependent CDM exhibits the best performance.

### 4.3 Stratifying across ranks

It has been generically noted [10, 12, 16] that the criteria individuals use for selecting top-ranked alternatives differs from those for lower-ranked alternatives, either due to decision fatigue or a true preference shift. As such, we consider the possibility of within-agent preference shift by stratifying the choice model by rank, and learning independent choice models at each rank position. A possible concern with this approach is that we end up with considerably less data for each model in this stratification, compared to estimating a single common model. To address this concern, we encourage models at neighboring ranks to be close to one another via Laplacian regularization, resulting in a **Laplacian-regularized stratified model** [37]. The methods of Laplacian-regularized stratification are closely related to popular methods for smoothing ( $\ell_2$ ) and trend filtering ( $\ell_1$ ) in temporal [17] and general graphical [32, 40] domains, where the underlying idea of parameter fusion dates back to at least the work of Land and Friedman [19, 33].

The stratification builds upon a base choice model—in this work, one of the three models summarized in Table 2. Taking the number of strata to be  $K$ , a stratified choice model is then the composition of  $K$  sub-models with parameters  $\theta = \{\theta_1, \dots, \theta_K\} \in \mathbb{R}^{K \times N}$ , where  $N$  is the number of parameters in the chosen base model.

The  $K$  models are regularized towards each other as dictated by an accompanying *regularization graph* [37]. In our case, rank-based stratification lends itself well to a common “path graph” for regularization, where models of adjacent ranks are connected by edges and thus regularized towards each other. Laplacian regularization here is then defined as:

$$r_{\mathcal{L}}(\theta) = \lambda_{\mathcal{L}} \sum_{i=2}^K \|\theta_i - \theta_{i-1}\|_2^2,$$

where  $\lambda_{\mathcal{L}}$  is a chosen Laplacian regularization strength, and  $r_{\mathcal{L}}$  is convex in  $\theta$ . Compared to the non-stratified objective in Eq. (3), the regularized, stratified objective function is the sum of  $K$  decoupled model losses (each with a local  $\ell_2$  regularization) and the Laplacian

regularization term:

$$F(D|\theta) = \sum_{k=1}^K [\ell(D_k; \theta_k) + r(\theta_k)] + r_{\mathcal{L}}(\theta). \quad (8)$$

Regularized stratified models feature two additional hyperparameters over their base models, the number of strata  $K$  and the Laplacian regularization gain  $\lambda_{\mathcal{L}}$ ; see Section 5 for a discussion of hyperparameter tuning.

## 5 MODEL SELECTION AND OPTIMIZATION

We briefly discuss the identifiability of the presented models, alongside details about hyperparameter tuning and optimization. A model is *identifiable* if no two distinct sets of parameters,  $\theta$  and  $\theta'$ , produce the same probability distributions over all choice sets  $S \in \mathcal{S}$ . Identifiability is crucial in settings where decisions are made based on interpreting parameter estimates. If the goal is solely to make decisions based on the resulting distributions only, e.g., from predictions or simulations, identifiability is not strictly necessary.

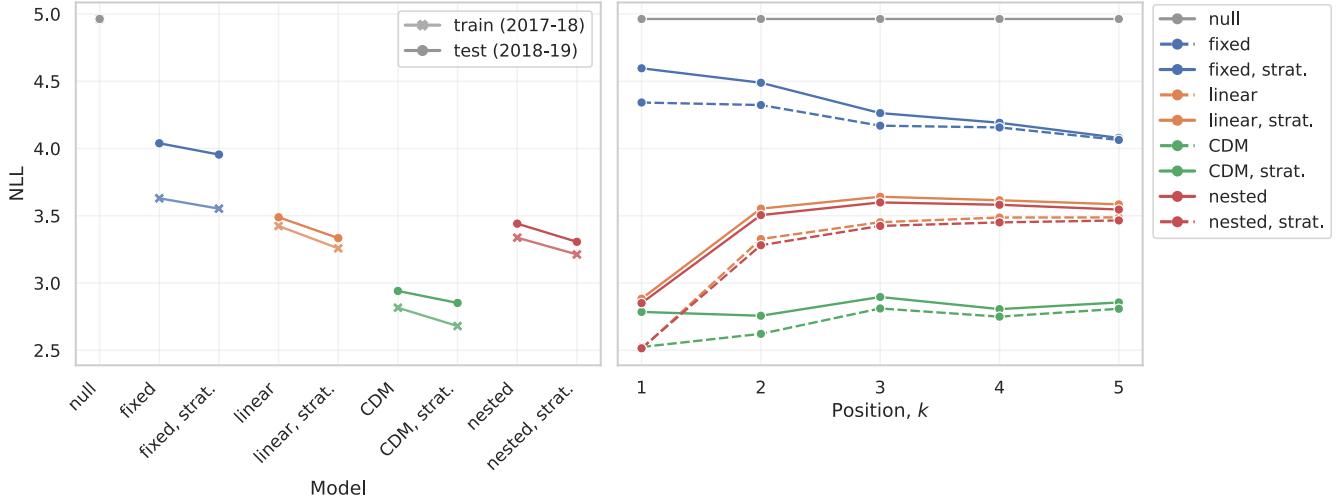
The traditional MNL family of ranking distributions are non-identifiable due to their shift-invariance. In this case, strategies for achieving identifiability are to fix one of the parameters, constrain their sum, or to apply regularization and obtain the minimum-norm parameter estimates [41]. In our work, we employ the latter strategy for the MNL and all other models, applying non-zero  $\ell_2$  regularization,  $r(\theta)$ , in the objective function and achieving identifiability by obtaining the minimum norm solution.

The models in this work introduce additional hyperparameters; the low-rank CDM requires the selection of the embedding dimension  $r$ , and a stratified model is specified by  $K$  and  $\lambda_{\mathcal{L}}$ ; the number of strata and amount of Laplacian regularization, respectively. We tune these hyperparameters via 5-fold cross validation within our training dataset, selecting the values that minimize validation loss. Figures illustrating our search over these hyperparameters are found in online Appendix B, with Table 3 summarizing the chosen values. With hyperparameter values selected and regularization in place, the models are fully specified and we proceed to train our models on the full 2017–18 dataset for testing on the 2018–19 dataset.

We run Adam [18], implemented in PyTorch, with default parameters,  $\text{lr} = 0.001$ ,  $\beta = (0.9, 0.999)$ ,  $\epsilon = 1e-8$ , adding  $\ell_2$  regularization with weight  $\lambda = 1e-5$  in accordance with our hyperparameter selection for  $\lambda$ . Model parameters are updated over batches of training data until reaching  $\text{max\_epoch} = 1000$  or convergence, i.e., when the absolute difference in losses is less than  $\epsilon = 1e-4$ . See online Appendix C for a discussion on the learned model parameters.

## 6 RESULTS

In this section, we evaluate and examine eight models—non-stratified and stratified versions of the fixed-effect MNL, linear MNL, CDM, and nested MNL models—trained on 2017–18 preference data and evaluated out-of-sample on 2018–19 data. We observe unique advantages of the context effects modeled by the CDM when benchmarked against the other models, and find that stratifying any model results in strictly (but marginally) better predictions, mostly for top (first) choices.



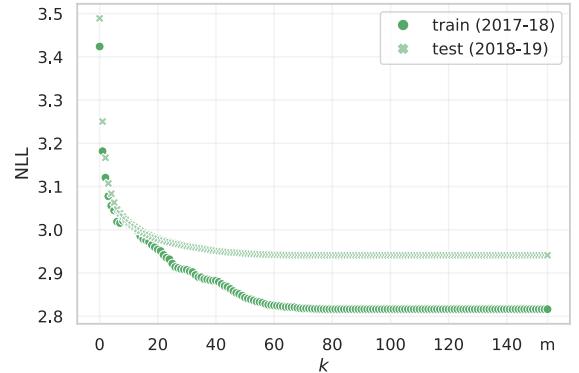
**Figure 1: Model negative log-likelihoods.** Left plot shows overall train and test losses (lower is better) on train (2017-18) and test (2018-19) datasets, respectively. Right plot shows out-of-sample test (2018-19) losses disaggregated by rank. CDM models display lowest overall test losses overall and down rank, and stratified models fit top-choices best.

## 6.1 Goodness of fit

Figure 1 depicts train and test negative log likelihood (NLL) losses on the left, and test losses disaggregated by rank on the right. We include a “null” model in the plots, representing uniform choices over programs, as a baseline reference point. We see that the CDM models, stratified or not, result in considerably lower test losses than the fixed-effect, linear, and nested MNL models overall. Stratifying provides modest decreases in overall test loss across all models.

On top choices, many families have priority access to one school in the district (e.g., sibling or PreK/TK priorities), and in most cases, rank these schools first. The linear and CDM models incorporate these priorities into the model, and therefore model top choices better than the fixed-effect model. The CDM leverages no additional information in the first choice as the context set is empty (i.e., no choices have been made). As such, there is negligible difference between the linear MNL and CDM models at position 1. However, after the relatively easy task of predicting top-choices, the CDM is able to leverage the choices made and separates itself from the lower-fidelity models. Stratifying yields a lower test loss for top choice across all three models, but quickly loses its advantage at lower ranks, likely due to diminishing training data at those positions (Cf. Table 1, households rank fewer than 10 programs on average).

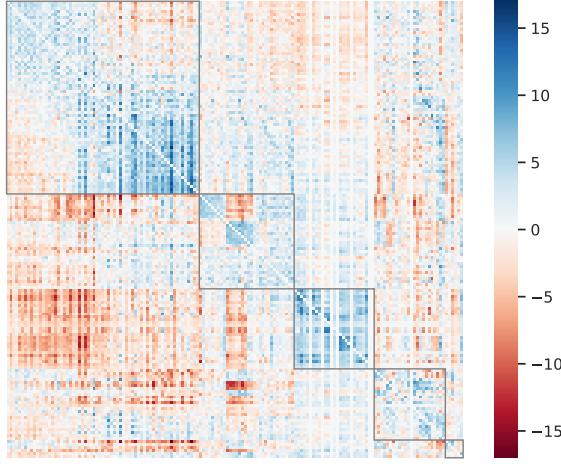
**Truncated top- $k$ -dependent CDM.** Recall from Eq. (6) that the CDM utility differs from the linear MNL model via a sum of pairwise interactions between alternatives and a context set,  $A$ . Throughout this work, the context set is taken to be the set of all previously-chosen alternatives, which has a powerful equivalence in expressivity (Theorems 1 and 2) to the standard CDM. As a robustness check, it is reasonable to ask which prior choices are *most relevant* to the context set. We evaluate several variations on the CDM model with the context being defined as the set of  $k$  top alternatives. Specifically, the utility is given by Eq. (6) with  $A_{ij} = \{r_{i1}, \dots, r_{il}\}$  for agent  $i$ ’s



**Figure 2: Negative log likelihoods of truncated top- $k$ -dependent CDMs where the context set is only the top- $k$  chosen alternatives. Here  $k = 0$  is equivalent to the linear MNL, and  $k = m$ , the total number of offered programs, is equivalent to the standard CDM.**

$j$ -th choice, where  $l = \min(j, k)$ . In other words, for choices made after position  $k$ , only the first  $k$  choices constitute the context.

Figure 2 presents the losses for these top- $k$ -dependent CDM models. When  $k = 0$ , the context set is always empty and the model is equivalent to the linear MNL. When  $k = m$ , the number of offered programs, we recover the backwards-dependent CDM considered everywhere else in this work. We find that even a minimal context set, e.g., the top-choice only ( $k = 1$ ), provides considerable improvement compared to the no-context linear model. That is, the information of what an agent chose first supplies the model with meaningful signal in making all down-rank predictions. That said,



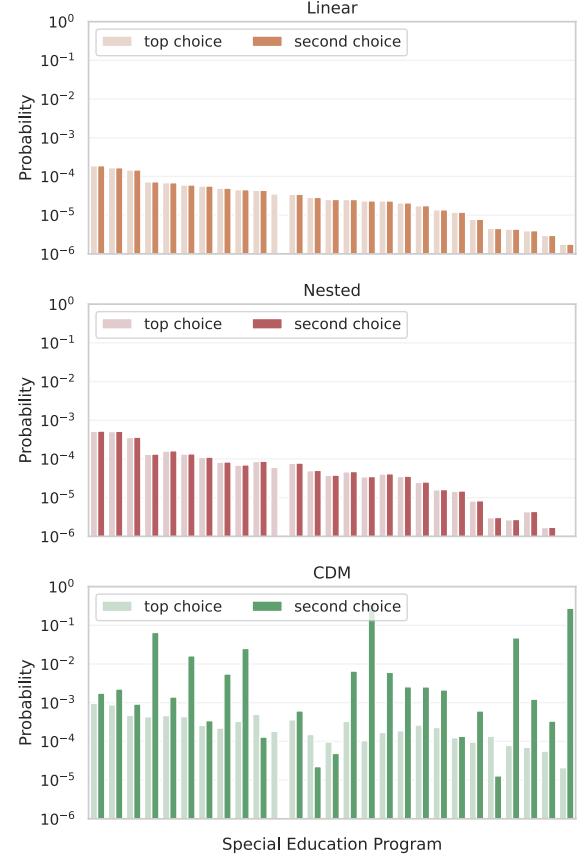
**Figure 3: The context effect matrix,  $U = TC^T$ , of the CDM.** Element  $u_{ij}$  is the utility boost program  $j$  receives from program  $i$  being in the context set. The block diagonal structure, highlighted with grey outlines, suggests that the CDM primarily (but not only) learns effects between like program types.

letting the context effect be linear in the full set of prior choices ( $k = m$ ) has measurable advantages.

*Interpreting context effects.* In Figure 3, we show the pairwise interactions  $U = TC^T$  estimated for the (non-stratified) backwards-dependent CDM. Element  $u_{ij} = t_j^T c_i$  is the utility boost that program  $j$  receives from chosen program  $i$  being in the context set. In the heatmap, programs on the  $x$ - and  $y$ -axes were arranged first by program type and then by (descending) popularity within each. From top/left to bottom/right, the program types are General Education (65), Spanish Language (32), Special Education (27), Chinese Language (24), and Miscellaneous Language (6) programs.

We see significant block structure in the matrix, suggesting that the CDM primarily (but not only) uses the context set to learn program type affinities. For example, the third block along the diagonal corresponds to Special Education programs, where we see a strong positive context effect. That is, once a family has ranked a special education program, it becomes much more likely that the family will rank other special education programs. This model behavior is highly intuitive, and is also beyond the behavior of an MNL model or any other model assuming independence. Put simply, the CDM’s use of context effects enables it to pick up on household signals, from the second choice and onward, that are otherwise not available *a priori* at the household level.

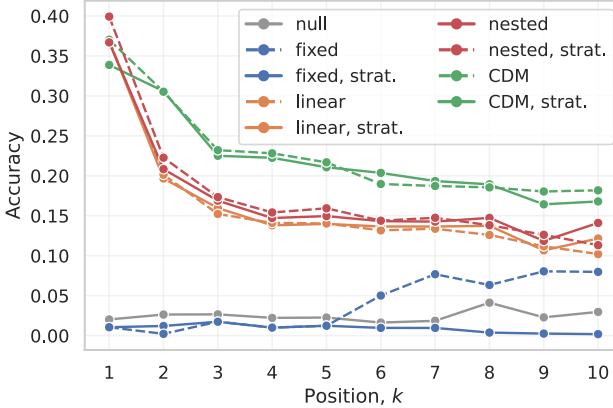
The block structure of  $U$  may seem to suggest good performance from a nested MNL model, as the latter explicitly clusters similar programs. Instead, in Figure 1 we find that the nested MNL shows only marginal gains over the linear MNL model and is not competitive with the CDM. This result sounds surprising, but is fairly intuitive; in models obeying IIA (such as the fixed-effect and linear MNL models), when an item is removed from the choice set, that item’s probability is proportionally redistributed to the remaining



**Figure 4: Top- and second-choice probabilities of the non-stratified linear, nested, and CDM models over special education programs for a sample household whose first choice was a special education program. The CDM model learns a drastically updated second-choice probability given the context of the top choice.**

alternatives for follow-up choices. The nested model instead allows the removed alternative’s probability to be *non-proportionally* distributed to the remaining items, specifically by favoring the alternatives in its nest (see online Appendix C.3 for details). However, in this setting, the choice universe and nests are relatively large, so the impact of redistributing already-small choice probabilities is marginal.

To illustrate how choice probabilities are redistributed in different models, Figure 4 showcases first- and second-choice probabilities by the non-stratified linear MNL, nested MNL, and CDM models over the special education subset for an example household in the district *who first chose a special education program*. We see that the CDM drastically alters its second-choice distribution, (correctly) boosting the likelihood of this household choosing another special education program, while the nested model’s top- and second-choice distributions are almost indistinguishable. Special education programs are low-probability selections in the data at large, and the nested paradigm can only marginally influence future



**Figure 5: Accuracy in  $k$ th Prediction.** The CDM model makes use of the provided context and generates the most accurate predictions at lower rank positions.

predictions when one such item is removed from the choice set. The CDM has a far greater ability to update its future distributions in the context of rare chosen items.

## 6.2 Down-rank prediction accuracy

Beyond in- and out-of-sample goodness of fit, we now consider the prediction quality of the models on the test dataset. Specifically, we task the models with making a prediction at rank position  $k$ , conditional on the first  $k-1$  choices made, resulting in an “accuracy in  $k$ th Prediction” evaluation metric. Recall that  $R_i$  denotes the true ranking of household  $i$ , where  $r_{ij}$  is the  $j$ th item in the ranking for  $j \leq k$ . Let  $R_{ik}$  be the set of their true top- $k$  choices,  $R_{ik} = \{r_{i1}, r_{i2}, \dots, r_{ik}\}$ . Denote by  $I_k = \{i : k_i \geq k\}$  the set of households who have ranked at least  $k$  alternatives. Then, given a choice model, denote by  $y_{ik}$  the *modal* prediction by the model at position  $k$ , i.e., the highest probability alternative over remaining programs, in the context of the previous  $k-1$  choices,  $A_{ik} = R_{i,k-1}$ ,

$$y_{ik} = \arg \max_{j \in S_{ik}} V_\theta(j|i, A_{ik}, S_{ik}),$$

where the representative utilities  $V(\cdot)$  are defined in Section 4. The metric is then given by

$$\text{Accuracy in } k\text{th Prediction} = \frac{1}{|I_k|} \sum_{i \in I_k} \mathbb{1}(y_{ik} = r_{ik}).$$

Figure 5 summarizes model performances on this metric. The CDM models are significantly more accurate in making down-rank predictions when given earlier choices, which is precisely the use-case of the contextual model. Stratification leads to improvements in down-rank predictions made by the fixed-effect MNL model, but has limited effect on the linear, nested and CDM models. It appears to learn that if a household has not already ranked the most popular programs, they won't be adding them later, as seen in Figure 8 of online Appendix C. Doing so, it outperforms its non-stratified counterpart beyond position 5.

We can also disaggregate these accuracies by sub-populations of interest, see online Appendix D. We find that the groups receiving

sibling and PreK/TK priorities have top choices that are relatively easy to predict, as their preferences are concentrated on their (typically singular) priority schools. All models generally under-perform on CTIP1, Hispanic/Latino, and Black student populations, relative to the broader population, for one of two reasons: either the subgroups demonstrate more varied preferences than other subgroups, or the training data was relatively small. Lastly, we see in Figure 12 that the CDM specializes in predicting down-rank choices for households with non-mainstream initial preferences.

## 7 CONCLUSION

In this work, we introduce rank-heterogeneous preference modeling for school choice and present two strategies, discrete choice context effects via a backwards-dependent CDM, and model stratification by rank position. Rank-heterogeneous models have the potential to leverage already-chosen alternatives when making down-rank predictions, or to broadly capture evolving household values down a ranking. We define and evaluate several metrics, finding that incorporating context terms in the utility dramatically decreases test loss over the linear and nested MNL models, capturing signals not present in covariates alone while also seeing particular improvements in modeling rare choices. The contextual model also generates more accurate predictions for list-completion tasks. Stratifying by rank yields improvements in top-choice accuracy across all models, but otherwise does not result in significant improvements or additional predictive power down-rank.

While rank-heterogeneous models enable school choice researchers to improve predictions and perform counterfactual analysis, our methods do not come without limitations. For one, the increased parametric complexity of the CDM and regularized stratification strategies raises, albeit mildly, model training times and data requirements relative to the MNL. Recent developments to the CDM [34] mitigates this problem by leveraging the model's block structure and learning interactions between program *attributes* rather than the programs themselves. Applying this work to the school choice setting would reduce complexity while uncovering context effects, a promising direction for future work. Another limitation stems from our model failing to generalize to new program offerings with undefined fixed-effects. Here, applying strategies for out-of-distribution prediction—such as establishing a prior on the fixed-effects of new program offerings based on the values for similar offerings—provide further directions for future work. Despite these limitations, we strongly encourage school choice researchers to consider rank-heterogeneous models in their preference modeling tasks for improved down-rank and rare-event prediction.

**Reproducibility.** The SFUSD data used in this work is not public, but implementations of all models as well as notebooks used to generate plots in this paper are available at: <https://github.com/ameloa/rankingmodels>.

## ACKNOWLEDGMENTS

We thank the San Francisco Unified School District for providing us access to choice data, specifically thanking SFUSD representatives Joseph Monardo, Jennifer Lebarre, Lauren Koehler, and Reed Levitt for helpful discussions. This work was supported in part by a gift from the Koret Foundation.

## REFERENCES

- [1] Atila Abdulkadiroğlu, Parag A. Pathak, Jonathan Schellenberg, and Christopher R. Walters. 2020. Do Parents Value School Effectiveness? *American Economic Review* 110, 5 (May 2020), 1502–39.
- [2] Nikhil Agarwal and Paulo Somaini. 2018. Demand Analysis Using Strategic Reports: An Application to a School Choice Mechanism. *Econometrica* 86, 2 (2018), 391–444.
- [3] Nikhil Agarwal and Paulo Somaini. 2020. Revealed Preference Analysis of School Choice Models. *Annual Review of Economics* 12, 1 (2020), 471–501.
- [4] Paul D. Allison and Nicholas A. Christakis. 1994. Logit Models for Sets of Ranked Items. *Sociological Methodology* 24 (1994), 199–228.
- [5] Itai Ashlagi and Afshin Nikzad. 2020. What matters in school choice tie-breaking? How competition guides design. *Journal of Economic Theory* 190 (2020), 105120.
- [6] Itai Ashlagi, Afshin Nikzad, and Assaf Romm. 2019. Assigning more students to their top choices: A comparison of tie-breaking rules. *Games and Economic Behavior* 115 (2019), 167–187.
- [7] Richard R. Batsell and John C. Polking. 1985. A New Class of Market Share Models. *Marketing Science* 4, 3 (1985), 177–198.
- [8] Amanda Bower and Laura Balzano. 2020. Preference modeling with context-dependent salient features. In *International Conference on Machine Learning*. PMLR, 1067–1077.
- [9] Caterina Calsamiglia, Chao Fu, and Maia Güell. 2020. Structural Estimation of a Model of School Choices: The Boston Mechanism versus Its Alternatives. *Journal of Political Economy* 128, 2 (2020), 642–680.
- [10] Randall G. Chapman and Richard Staelin. 1982. Exploiting Rank Ordered Choice Set Data within the Stochastic Utility Model. *Journal of Marketing Research* 19, 3 (1982), 288–301.
- [11] Douglas E. Critchlow, Michael A. Fligner, and Joseph S. Verducci. 1991. Probability models on rankings. *Journal of Mathematical Psychology* 35, 3 (1991), 294–318.
- [12] Dennis Fok, Richard Paap, and Bram Van Dijk. 2012. A rank-ordered logit model with unobserved heterogeneity in ranking capabilities. *Journal of Applied Econometrics* 27, 5 (2012), 831–846.
- [13] D. Gale and L. S. Shapley. 1962. College Admissions and the Stability of Marriage. *The American Mathematical Monthly* 69, 1 (1962), 9–15.
- [14] Monique De Haan, Pieter Gautier, Hessel Oosterbeek, and Bas Van der Klaauw. 2015. The Performance of School Assignment Mechanisms in Practice. *Journal of Political Economy* (2015), null.
- [15] Justine S. Hastings, Thomas J. Kane, and Douglas O. Staiger. 2008. Heterogeneous Preferences and the Efficacy of Public School Choice.
- [16] Jerry A. Hausman and Paul A. Ruud. 1987. Specifying and testing econometric models for rank-ordered data. *Journal of Econometrics* 34, 1 (1987), 83–104.
- [17] Seung-Jean Kim, Kwangmoo Koh, Stephen Boyd, and Dmitry Gorinevsky. 2009.  $\| \cdot \|_1$  trend filtering. *SIAM review* 51, 2 (2009), 339–360.
- [18] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization.
- [19] Stephanie R Land and Jerome H Friedman. 1997. Variable fusion: A new adaptive signal regression method.
- [20] Mariana Laverde. 2022. *Distance to Schools and Equal Access in School Choice Systems*. Working Papers 2022-002. Human Capital and Economic Opportunity Working Group.
- [21] R. Duncan Luce. 1959. *Individual Choice Behavior: A Theoretical analysis*. Wiley, New York, NY, USA.
- [22] R. Duncan Luce. 1977. The choice axiom after twenty years. *Journal of Mathematical Psychology* 15, 3 (1977), 215–233.
- [23] Daniel McFadden. 1978. Modelling the choice of residential location. *Spatial Interaction Theory and Planning Models* (1978).
- [24] Daniel McFadden and Kenneth Train. 2000. Mixed MNL models for discrete response. *Journal of Applied Econometrics* 15, 5 (2000), 447–470.
- [25] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space.
- [26] Parag A. Pathak and Peng Shi. 2014. Demand Modeling, Forecasting, and Counterfactuals, Part I.
- [27] Parag A Pathak and Peng Shi. 2017. *How Well Do Structural Demand Models Work? Counterfactual Predictions in School Choice*. Working Paper 24017. National Bureau of Economic Research.
- [28] R. L. Plackett. 1968. Random Permutations. *Journal of the Royal Statistical Society: Series B (Methodological)* 30, 3 (1968), 517–534.
- [29] Girish N. Punj and Richard Staelin. 1983. A Model of Consumer Information Search Behavior for New Automobiles. *Journal of Consumer Research* 9, 4 (03 1983), 366–380.
- [30] Arjun Seshadri, Alexander Peysakhovich, and Johan Ugander. 2019. Discovering Context Effects from Raw Choice Data. *CoRR* abs/1902.03266 (2019). arXiv:1902.03266
- [31] Arjun Seshadri, Stephen Ragaini, and Johan Ugander. 2020. Learning Rich Rankings. In *Advances in Neural Information Processing Systems*, Vol. 33. 9435–9446.
- [32] Alexander J Smola and Risi Kondor. 2003. Kernels and regularization on graphs. In *Learning theory and kernel machines*. Springer, 144–158.
- [33] Robert Tibshirani, Michael Saunders, Saharon Rosset, Ji Zhu, and Keith Knight. 2005. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67, 1 (2005), 91–108.
- [34] Kiran Tomlinson and Austin R Benson. 2021. Learning interpretable feature context effects in discrete choice. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 1582–1592.
- [35] Kenneth Train. 2009. *Discrete Choice Methods with Simulation* (second edition ed.). Cambridge University Press, New York, NY, USA.
- [36] Jennifer S Trueblood, Scott D Brown, Andrew Heathcote, and Jerome R Busemeyer. 2013. Not just for consumers: Context effects are fundamental to decision making. *Psychological science* 24, 6 (2013), 901–908.
- [37] Jonathan Tuck, Shane Barratt, and Stephen Boyd. 2021. A Distributed Method for Fitting Laplacian Regularized Stratified Models. *J. Mach. Learn. Res.* 22, 1, Article 60 (Jan 2021), 37 pages.
- [38] Amos Tversky. 1969. Intransitivity of Preferences. *Psychological Review* 76, 1 (1969), 31–48.
- [39] Amos Tversky and Itamar Simonson. 1993. Context-Dependent Preferences. *Management Science* 39, 10 (1993), 1179–1189.
- [40] Yu-Xiang Wang, James Sharpnack, Alex Smola, and Ryan Tibshirani. 2015. Trend filtering on graphs. In *Artificial Intelligence and Statistics*. PMLR, 1042–1050.
- [41] Lirong Xia. 2019. Learning and Decision-Making from Rank Data. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 13 (02 2019), 1–159. <https://doi.org/10.2200/S00876ED1V01Y201810AIM040>