A Game-theoretic Analysis of Catalog Optimization

Joel Oren, Nina Narodytska and Craig Boutilier

Abstract

Vendors of all types face the problem of selecting a slate of product offerings—their *assortment* or *catalog*—that will maximize their profits. The profitability of a catalog is determined by both customer preferences and the offerings of their competitors. We develop a game-theoretic model for analyzing the vendor *catalog optimization* problem in the face of competing vendors. We show that computing a best response is intractable in general, but can be solved by dynamic programming given certain informational or structural assumptions about consumer preferences. We also analyze conditions under which pure Nash equilibria exist and provide several price of anarchy/stability results

1 Introduction

Vendors of retail products and services typically plan their offerings to maximize the revenue/profits obtainable from a (predicted or actual) customer population. In many cases, the prices of these items are fixed or strongly suggested exogenously (e.g., a vendor opening a new branch of a retail chain in a mall). The problem of optimizing the collection of products offered or presented to customers is often subtle: highly profitable products may appeal only to a small subset of customers, while offering lower-value products that appeal to a larger market may undercut both profits and sales of the higher-value offerings. This problem is known as optimization of *assortment* (deciding which products to stock), or *catalogs* (which products to promote in a catalog, web site, etc.) and is faced by traditional (offline and online) retailers as well as multi-seller platforms like Amazon or Ebay.

Complicating the picture is the presence of *competing vendors*. A target customer may choose to purchase from a vendor's competitor if the competitor offers a more preferred product. Thus the selection of the revenue maximizing catalog also depends on the offerings of one's competitors. This is naturally formulated as a game, the *competitive catalog selection game*. In this paper, we formulate and analyze various aspects of this game. Roughly speaking, the model assumes a collection of strategic vendors, each of whom can select a *catalog*, i.e., a subset of some underlying collection of products within a specific category, for sale to a target audience or market of unit-demand customers. All prices are fixed exogenously and are beyond the control of the vendors. Each customer has preferences over products (which can depend on the prices), and purchases her most preferred from the set of all offered products.

We analyze several key properties of this game under a variety of conditions w.r.t. the structure of, and vendor information about, consumer preferences. We consider two conceptually distinct models of information. In the *complete information* model, vendors know the true consumer preference profile, i.e., the precise ranking of each consumer for all products. In the *partial information* model, vendors have (common, prior) *probabilistic beliefs* over profiles and must maximize expected revenue. We first consider the algorithmic task of computing a vendor's best response, i.e., the optimal catalog given the catalogs of her competitors. We show that this is hard to compute and to approximate in the complete information model. However, we provide an efficient dynamic programming (DP) method for the partial information model when preferences are drawn i.i.d. from the (widely used) Mallows model (and mixtures thereof). In the special case of uniformly random preferences, or *impartial culture*, DP reduces to a simple greedy algorithm. We also briefly describe a special case of *single-peaked truncated preferences* that admits a DP algorithm under full information.

We then analyze the stability of the game. We describe (straightforward) instances of the complete information catalog game where no pure Nash equilibrium exists, and show this instability persists even if vendors are restricted to small sets of items (even singletons). In contrast, under impartial culture, we show that pure equilibra exists using simple best-response dynamics (and can be computed efficiently). Finally, we provide several *Price of Anarchy/Stability* results, showing that vendor welfare in the best/worst pure equilibrium in a partial information game may be linear the total number of products, and provide additional analysis of special cases in which all vendors have identical product sets.

2 Preliminaries

We consider a game with k strategic vendors, each vendor j having access to a set $C^j = \{c_{j1}, \ldots, c_{jm_j}\}$ of m_j items. We do not require disjointness of these sets, though we sometimes assume this in some of the exposition below (w.l.o.g.). Let $C = \bigcup_{j=1,\ldots,k} C^j$, and $m = \sum_{j=1}^{k} m_j$. Let $\mathcal{L}_m$ denote the set of all possible ordinal preference rankings over (or permutations of) C. We assume that each item $c \in C$ has an *exogenously fixed*, bounded price $p(c) \geq 0$. Let $\mathbf{p}$ denote the price vector over C. W.l.o.g., we assume $\mathbf{p} \in [0,1]^m$. Each vendor offers a *catalog* $R^j \subseteq C^j$ from which consumers can purchase items. Let $\mathbf{R} = (R^1, \ldots, R^k)$.

We assume a set of *unit-demand* consumers $N = \{1, \ldots, n\}$. Each consumer i has a *strict preference ordering* $\sigma_i \in \mathcal{L}_m$, representing her preferences over items C. If $\sigma_i(c) < \sigma_i(c')$, then i prefers c to c' (given their fixed prices); i.e. i *ranks* c above c'. While we focus on ordinal preferences, σ_i can be a ranking induced from i's intrinsic valuation for items and the prices. Let $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$ be the *preference profile*. We could generalize the model to allow consumers to "truncate" their preferences (e.g., consider certain products unacceptable or too expensive) and to allow ties/indifference; we do not do so for ease of exposition, but our results here can be generalized as appropriate.

We assume a simplified supply/demand model: each vendor has unlimited inventory of any product she offers and no production cost.[1] Given a (non-empty) set of offered items $A = \bigcup_{j=1}^{k} R^j$, each $i \in N$ buys their most preferred item in A, i.e., $top_i(A) = \arg\min_{c \in A} \sigma_i(c)$. Let $top_i(\mathbf{R}) = top_i(\bigcup_{j=1}^{k} R^j)$. This determines a (single-shot) game $\mathcal{G} = (C^1, \ldots, C^k, \mathbf{p}, N, \boldsymbol{\sigma})$, in which vendor j's strategies are catalogs $R_j \subseteq C^j$, and her revenue (payoff) is determined by the *strategy profile* $\mathbf{R}$. In models where the C^j are not disjoint, if consumer i selects a c that occurs in catalogs of several vendors, we assume $p(c)$ is split evenly among them (as if i randomized her purchase). Let $o_c(\mathbf{R}) = |\{j \mid c \in R^j, j \leq k\}|$ be the number of occurrences of item c across all catalogs. The *revenue/payoff* of vendor j is:

$$r_j(\mathbf{R}) = r_j(R^j, \mathbf{R}^{-\mathbf{j}}) = \sum_{i \in N : top_i(\mathbf{R}) \in R^j} \frac{p(top_i(\mathbf{R}))}{o_{top_i(\mathbf{R})}(\mathbf{R})}.$$

Vendor j's *best response* to $\mathbf{R}^{-\mathbf{j}}$ is the subset $R^j \subseteq C^j$ that maximizes her revenue, given catalogs of the other vendors.

Probabilistic preference models. In the *full information* game, we assume vendors have full knowledge of the consumer preference profile $\boldsymbol{\sigma}$. In the *partial information* setting, vendors instead have only a common prior belief, or a *distribution* over preference profiles. To model this scenario, we assume that consumer preferences are drawn i.i.d. from a *Mallows φ-distribution* (or a mixture thereof), a probabilistic model of rankings widely used in statistics, machine learning, econometrics and social choice [11, 12]. The Mallows model is specified by two parameters, a *reference ranking* $\hat{\pi} \in \mathcal{L}_m$ and a *dispersion parameter* φ (controlling variance). The probability of a ranking σ in this

[1] Prices can also be considered to reflect net revenue, so this is w.l.o.g. if costs are fixed and per-unit.

model is $\Pr(\sigma) = \varphi^{\tau(\sigma,\hat{\pi})}/T_m$, where τ is the Kendall-tau (or swap) distance between two rankings, and T_m is a normalization constant. It is well-known that T_m is equal to $T_m = \prod_{t=1}^{m} Z_t$, where $Z_t = \sum_{d=1}^{t} \varphi^{d-1}$. When $\varphi = 1$, one obtains the uniform distribution over $\mathcal{L}_m$, or the *impartial culture (IC)* model, widely used in social choice.

Related Work. Work on assortment optimization is prevalent in management science, and several models bear tight connection to ours. Some models consider *non-strategic* optimization on the part of single vendors (e.g., [15]), where the aim is to select a revenue-maximizing catalog assuming some consumer preference model. More directly relevant is work addressing the *strategic aspects* of this problem in the face of competition, assessing both computation and existence of equilibria. [13] study the efficiency of equilibria of vendors competing for *limited* shelf space, where key differences exist from our approach in their consumer preference model (as well as constraints on shelf space). They address both exogenous and endogenous prices. Li et al. [9] investigate similar models. Hohnon et al. [7] investigate (non-competitive) variants of the problem using a rank-based preference model similar to ours.

Our model can be thought of an extension *multi-winner social choice (MWCS)*. In MWSC, the goal is to select a "slate" of items given a set of agent preferences, and has application to legislature/committee selection [2, 14], facility location, and group (e.g., consumer) decision-making [8, 10, 16]. In our setting, we have *multiple decision makers* rather than a single social choice mechanism, and each of these are strategic. Many of the computational tasks in MWSC can be implemented as extensions of our model (e.g., imposing different combinatorial restrictions on the strategies, considering various classes of preference models). Our model is somewhat related to the task of optimal price-setting mechanisms for auctioning items in unlimited supply to unit-demand bidders [5].

3 Best Responses under Full Information

We begin with the task of computing a vendor best response to a competitors' strategy profile, i.e., given profile $\mathbf{R^{-j}}$, finding the catalog R^j that optimizes j's payoff. In the non-competitive version of this problem, in which no other vendors offer products, j's optimal catalog is trivial: she should offer only her product with the maximal price.[2] Also notice that if all of her prices are identical, trivially she should offer her entire set C^j. In general, however, best-response computation is hard to approximate beyond a constant factor.

Theorem 1. *Computing a best response is Max-SNP hard.*

Proof. We provide an approximation-preserving reduction from 3SAT-5, which is known to be Max-SNP hard [4]. Take as input m DNF clauses $(\varphi_1, \ldots, \varphi_m)$ over n variables $(x_1, \ldots, x_n)$.

The reduction: Create a set of items C^j with two items for each variable: a_i, b_i, corresponding to a $True$ assignment to either x_i or its negation, plus one auxiliary item t. Set the price of all items in $\{a_i, b_i\}_{i=1,\ldots,n}$ to 1 and of t to 1.5. Create two sets of consumers:
Set 1: Validity rankings. These consumers encode validity constraints on assignments to $x_1, \ldots, x_n$, i.e., that exactly one of $\{a_i, b_i\}$ is $True$. For $i \in [n]$ create partial rankings:

$$\sigma_{i1} : a_i \succ t, \quad \sigma_{i2} : b_i \succ t, \quad \sigma_{i3} : a_i \succ b_i, \quad \sigma_{i4} : a_i \succ b_i$$

(The competitor items are always ranked below the items in C^j in all of the preferences). Taking t can never hurt, so we assume it is always chosen. If both a_i and b_i are selected, payoff is 4. If neither is selected, payoff is 3. If one is chosen, payoff is 4.5.

[2] This is not the case if preferences are truncated.

Set 2: Clause rankings. For each clause of the form $\varphi_j = \ell_{j_1} \vee \ell_{j_2} \vee \ell_{j_3}$, create a ranking: σ_j : $f(\ell_{j_1}) > f(\ell_{j_2}) > f(\ell_{j_3})$, where $f(\ell_{j_i}) = a_{j_i}$ if ℓ_{j_i} corresponds to variable x_{j_i} in its non-negated form, and b_{j_i} otherwise. If φ_j is satisfied, at least one item corresponding to ℓ_{j_1}, ℓ_{j_2}, and ℓ_{j_3} is selected, which gives an extra payoff of 1 from ranking σ_j.

Since 3SAT-5 is Max-SNP hard, there exists a constant $\epsilon > 0$ s.t. it is NP-hard to distinguish a satisfiable formula from a formula that is at most $(1-\epsilon)$-satisfiable. By definition, each variable is in exactly five 3CNF clauses, and so $m = 5n/3$. If φ is only $(1-\epsilon)$-satisfiable, the maximum value we can obtain is $(1-\epsilon)m + 4.5n = (1-\epsilon)m + 27m/10$, hence it is NP-hard to distinguish between cases with a profit of $m + 4.5n$ from cases in which only a $(1-\delta)$ fraction of $m + 4.5n$ profit, where $\delta = \epsilon/(1 + 27/10)$. □

The construction uses preferences of length at most 3, and item prices a factor of 1.5 from each other. Thus, selecting all items gives a 1.5-*approximation to the optimal catalog*. In general, if there is a constant $\beta > 1$ s.t. for every two distinct items $a, b \in C^j$, $\frac{p(a)}{\beta} \leq p(b) \leq \beta \cdot p(a)$, then selecting all items in C^j is a β-approximation to the optimal catalog.

The above hardness result suggests two directions for further investigation. The first is developing approximations, a topic we leave to future research. The second, is the study of the best-response problem under various restrictions on items, prices, or preferences. As an example, consider the case in which the agent preferences are *single-peaked* [1]. In the following section (Section 4), we show that with such preferences, and under certain other mild conditions, best-response computation is amenable to a tractable dynamic programming algorithm.

4 Single-Peaked, Bounded-Length Truncated Preferences

Given the hardness of the problem in the full information setting, we impose the following natural restriction on the preferences of the consumers, which will allow us to efficiently find the best response for a given vendor. Assume that the consumers are *single-peaked*, for which a formal definition will be given momentarily. The class of single-peaked preference has been heavily studied in social choice theory [1] as it is well-suited for modeling user preferences, like political preferences or distance-dependant preferences.

It is well-known that the single-peaked preferences assumption can make computational problems over preferences tractable even if the problem is intractable over arbitrary preferences (e.g., [3]).

In addition to the "single-peakedness" of the preferences, we impose one more restriction on the parameters of the game. To motivate it, consider a scenario in which many of the vendor sets consist of very few items. In any reasonable realization of the game, these vendors will resort to selecting all, or most of their items. If this behavior is sufficiently extensive, most of the consumers will have at least one product, that belong to one such "weak" vendor, ranked relatively high in her ranking. We are therefore interested in cases where this behavior results in consumers with effectively *short truncated* preferences; i.e., preferences that tend to have a high-ranking offered competitor item.[3]

We now formally define the notion of truncated, single-peaked preference profiles (we combine the above two assumptions in the definition, for succinctness)..

Definition 2. *For a vendor $j \in [k]$, and the set of strategies of the remaining vendors $\mathbf{R^{-j}}$, the preferences $(\pi_1, \ldots, \pi_n)$ are said to be single-peaked, L-truncated if:*

1. Bounded number of relevant items per consumer: *For every consumer $i \in N$, letting, $t_i = \arg\min_{c \in \bigcup_{j' \neq j} R^{j'}} \pi_i(c)$, we have that $t_i \leq L + 1$. For convenience, we let $S_i = \{c \in C^j : \pi_i(c) < t_i\}$; i.e., the items in C^j that are relevant to consumer i.*

[3] Note that the proof of Theorem 1 already made use of preferences of length at most 3.

2. Single-peaked preferences: *There exists an ordering $\hat{\pi} \in \mathcal{L}_m$, of the items in C (the* axis*), such that the following holds for each consumer $i \in N$. There exists a distinct item $c \in C$ (the* peak*), for which every two distinct candidates $c', c'' \in S_i \setminus \{c\}$, having $\hat{\pi}(c) > \hat{\pi}(c') > \pi(c'')$ or $\hat{\pi}(c) < \hat{\pi}(c') < \hat{\pi}(c'')$ implies $\pi_i(c') < \pi_i(c'')$.*

Note that we are only imposing the single-peaked structure over the preference prefixes that are relevant to vendor j, instead of on the entire preference profile.

We assume that for vendor j, his items are labeled according to their order along the axis: for two distinct items $c_{jt}, c_{j\ell} \in C^j$, $t < \ell$ iff $\hat{\pi}(c_{jt}) < \hat{\pi}(c_{j\ell})$. Also, recall that we can assume w.l.o.g. that no items in C^j are also offered by a competitor, as this would impose vendor j to include it as well in his strategy.

The following theorem argues that optimizing the best-response of a vendor is tractable, under our two assumptions.

Theorem 3. *Let $\mathbf{R}^{-\mathbf{j}}$ be the strategy profile of all vendors expect vendor j. Then in the case where the preference are single-peaked, L-truncated, there exists a polynomial-time algorithm that finds the best response for jth vendor, R^j, in $O(2^L m \cdot n)$ time.*

The following claim will be instrumental for efficiently finding a best response for the given vendor.

Claim 4. *Consider a consumer $i \in N$ with a single-peaked L-truncated preference π_i over items $S_i \subseteq C^j$. Then $\max_{c,c' \in C^j} |\hat{\pi}(c) - \hat{\pi}(c')| < L$.*

Proof. Suppose that there exists such a consumer $i \in N$ with an L-truncated preference σ_i over the subset $S_i \subseteq C^j$ that contains two items c', c'' such that $|\hat{\pi}(c') - \hat{\pi}(c'')| > L$.

Then by the pigeonhole principle there has to exist an item $c^* \in C \backslash S_i$ such that c^* lies between c' and c'' in the axis $\hat{\pi}$, which contradicts the fact that all the consumers have single-peaked preferences. □

Intuitively, Claim 4 implies that the decision of whether or not to take an item c cannot affect consumers with truncated preferences over sets of items that contain items at distance greater than L from c, w.r.t. $\hat{\pi}$.

We now describe the dynamic programming algorithm for optimizing the best-response (Algorithm 1. Letting $L \leq t \leq m$, suppose that for each subset S' of items in $\{\hat{\pi}^{-1}(t - L), \ldots, \hat{\pi}^{-1}(t-1)\}$ we computed the optimal slate that includes *all* of items in S' and some subset of $\{\hat{\pi}^{-1}(1), \ldots, \hat{\pi}^{-1}(t - L - 1)\}$. To compute an optimal value for analogous S, such that $S \subseteq \{\hat{\pi}^{-1}(t - L + 1), \ldots, \hat{\pi}^{-1}(t)\}$ we only need to consider at most two subsets: $S \setminus \{\hat{\pi}^{-1}(t)\}$, and $S \setminus \{\pi^{-1}(t)\} \cup \{\hat{\pi}^{-1}(t - L)\}$.

Note that the correctness of Algorithm 1 follows immediately from Claim 4. Indeed, the decision of whether or not to take item c_{jt} does not affect the revenue due to consumers with peaks at distance greater than L from c_{jt} on the axis $\hat{\pi}$, thus implying the appropriate optimal substructure property. We defer the full proof of correctness to a full version of the paper. Note that the running time of the algorithm is $O(2^L \cdot n \cdot m^2)$, which is polynomial in n and m if $L = O(\log m)$.

5 Best Responses under Partial Information

The full information model in which vendors know consumer preferences precisely is unrealistic in many settings. We now address best response computation in the *partial information game*, under several distinct forms of beliefs.

We consider a different model of the vendors' beliefs about the consumer preferences. We assume that the preference profiles are drawn from a probability distribution D that is fully known to

Algorithm 1: The dynamic programming algorithm for finding the optimal unconstrained slate for an single-peaked, L-truncated preference profile.

Input: A single-peaked, L-truncated preference profile $(\pi_1, \ldots, \pi_n) \in \mathcal{L}_m^n$. An underlying axis $\hat{\pi} \in \mathcal{L}_m$, The preferences profile $\mathbf{R^{-j}}$, composed of strategies of all vendors $j' \neq j$.

Notation: *For a binary vector $v \in \{0,1\}^L$, we let $B(v)$ denote the decimal representation of v: $B(v) = \sum_{t=1}^{k} 2^{t-1}$.*

For a subset $S \subseteq \{c_{j(t-L+1)}, \ldots, c_{jt}\}$ we let $\mathbb{1}_{S,t}$ denote the length-L characteristic vector of S w.r.t. $\{c_{j(t-L+1)}, \ldots, c_{jt}\}$. That is, $\mathbb{1}_{S,t}(d) = 1$ iff $c_{j(t-L+d)} \in S$ and 0 otherwise, for $d = 1, \ldots, L$.

Let M be an m_j by 2^L table.

For $t \in \{1, \ldots, m_j\}$ and $S \subseteq \{c_{j(t-L+1)}, \ldots, c_{jt}\}$, $M[t, B(\mathbb{1}_{S,t})]$ contains the optimal solution for the problem of optimizing the slate using items from $\{c_{j1}, \ldots, c_{jt}\}$, such that $M[t, B(\mathbb{1}_{S,t})] \setminus \{c_{j1}, \ldots, c_{j(t-L)}\} = S$.

for $t \to 1$ *to* L **do**
  foreach $S \subseteq \{c_{j1}, \ldots, c_{jt}\}$ **do**
    $M[t, B(\mathbb{1}_{S,t})] \leftarrow S$
for $t \leftarrow L+1$ *to* m_j **do**
  foreach $S \subseteq \{c_{j(t-L+1)}, \ldots, c_{jt}\}$ **do**
    $S_1 \leftarrow M[t-1, B(\mathbb{1}_{S \setminus \{c_{jt}\}, t-1})]$
    $S_2 \leftarrow M[i-1, B(\mathbb{1}_{S \setminus \{c_{jt}\} \cup \{c_{j(t-L)}\}, t-1})]$
    if $r_j(S_1 \cup S, \mathbf{R^{-j}}) \geq r_j(S_2 \cup S, \mathbf{R^{-j}})$ **then**
      $M[t, B(\mathbb{1}_{S,t})] \leftarrow S_1 \cup S$
    else
      $M[t, B(\mathbb{1}_{S,t})] \leftarrow S_2 \cup S$
return $\arg\max_{S \subseteq \{c_{j(m_j-L+1)}, \ldots, c_{jm_j}\}} M[m, B(\mathbb{1}_{S,m_j})]$

---

all of the vendors. In the induced game $\mathcal{G} = (C^1, \ldots, C^k, \mathbf{p}, N, D)$, player i's revenue for a strategy profile $(R^1, \ldots, R^k)$ is defined to be his expected revenue, when taking the expectation over the instantiation of the preference profile. In this section, we pursue three types of distributions: the Impartial Culture, the Mallows distribution, and a mixture of Mallows. Note that each distribution generalizes its preceding generalization.

While we do not require disjointness of vendor item sets, it is not hard to see that, for the purpose of selecting a best response, if a vendor's set contains an item that is currently offered by a competitor, that item must be included in the vendor's best response. Hence, for ease of exposition, in this section we assume that $C^1, \ldots, C^k$ are all disjoint.

Impartial culture. We begin with the case where consumer preferences are (believed to be) distributed according to IC, i.e., each consumer's preference is drawn i.i.d. from $\mathcal{L}_m$. Computing a best response for vendor j under IC is straightforward: assume competitor profile $\mathbf{R^{-j}}$, and let $\ell^j = \sum_{j' \neq j} |R^{j'}|$. We relabel the items in C^j so that $p(c_{j1}) \geq \cdots \geq p(c_{jm_j})$, and define the length t prefix of this item vector: $T_t = \{c_{j1}, \ldots, c_{jt}\}$ and $T_0 = \emptyset$.

For any catalog $R \subseteq C^j$, j's expected profit is $r_j(R, \mathbf{R^{-j}}) = \sum_{c \in R} p(c)/(\ell^j + |R|)$. Let $t^* = \arg\max_{1 < t \leq m_j} \{r_j(T_t, \mathbf{R^{-j}}) > r_j(T_{t-1}, \mathbf{R^{-j}})\}$. The best response is then T_{t^*}, i.e., we greedily add items in decreasing order of price as long as the expected revenue is increased by these additions. The addition of *any* items beyond that point cannot contribute to j's expected profit. The optimality of the greedy algorithm can be proven using an elementary exchange argument.

Mallows models. We now address the broader class of distributions, where preferences are drawn i.i.d. from a Mallows model (note that impartial culture is a special case). Best responses can be computed using a dynamic programming method. Let $\mathbf{R^{-j}}$ be a competitor strategy profile for vendor j, and beliefs given by Mallows model $(\hat{\pi}, \varphi)$. For convenience, we assume that $\hat{\pi}$ is restricted to elements of $\mathbf{R^{-j}}$. As above, let $\ell^j = \sum_{j' \neq j} |R^{j'}|$. We assume w.l.o.g. that items in C^j are ordered based on their ranks in $\hat{\pi}$.

For integer $s \leq m_j$ and index $t = s, \ldots, m_j$, Algorithm 2) recursively computes the optimal catalog of size s consisting of items from subset $\{c_{j1}, \ldots, c_{jt}\}$. Given $R^j \subseteq C^j$ and $c \in R^j$, let $\hat{\pi}_{R^j}(c)$ be the rank of c in the reduced ranking $\hat{\pi}_{R^j}$, obtained by deleting all items in $C^j \setminus R^j$. Using the recursive nature of the Mallows distribution to compute the revenue of vendor j. For $c_{jt} \in C^j$, and a set of previously selected items $S \subseteq \{c_{j1}, \ldots, c_{j(t-1)}\}$, s.t. $|S| = s-1$, the probability that c_{jt} is selected (i.e., ranked first) can be shown to be $\varphi^{\hat{\pi}_S(c_{jt})-1}/Z_{\ell^j+s}$, where Z_{ℓ^j+s} is the normalizing term. Thus, if the expected revenue of selecting S is $r_j(S, \mathbf{R^{-j}})$, the expected revenue of adding c_{jt} to S is: $r_j(S \cup \{c_{jt}\}, \mathbf{R^{-j}}) = (r_j(S, \mathbf{R^{-j}}) \cdot Z_{\ell^j+s-1} + \varphi^{\hat{\pi}_{S \cup \{c_{jt}\}}-1} \cdot p(c_{jt}))/Z_{\ell^j+s}$.

Algorithm 2: Dynamic programming algorithm for best-response given a Mallows distribution.

Assume $\ell^j = \sum_{j' \neq j} |R^{j'}|$, and for a non-negative integer q, let $\tau_q = Z_{\ell^j+q} = \sum_{d=0}^{\ell^j+q-1} \varphi^d$.
for $s \leftarrow 1$ *to* m_j **do**
  | Let $v_s, S_s \leftarrow$ `OptimizeSlate`(s, s)
Return S_s with maximal value v_s.

`OptimizeSlate` (s, t)
Input: A Mallows distribution with parameters $(\hat{\pi}, \varphi)$.
Output: Optimal slate $R^j \subseteq \{c_{j1}, \ldots, c_{jt}\}$, s.t. $|R^j| = s$.
if $s = 0$ **then**
  | **return** $0, \emptyset$
if $s = t$ **then**
  | $S \leftarrow \{c_{j1}, \ldots, c_{jt}\}$
  | $V \leftarrow \sum_{d=1}^{s} p(c_{jd}) \cdot \varphi^{\hat{\pi}_S(c_{jd})-1}/\tau_s$
  | **return** S, V
$v_0, S_0 \leftarrow$ `OptimizeSlate`$(s, t-1)$
$v_1, S_1 \leftarrow$ `OptimizeSlate`$(s-1, t-1)$
$S_1 \leftarrow S_1 \cup \{c_{jt}\}$
$v_1 \leftarrow (v_1 \cdot Z_{s-1} + p(c_{jt}) \cdot \varphi^{\hat{\pi}_{S_1}(c_{jt})-1})/\tau_s$
if $v_0 \geq v_1$ **then**
  | **return** v_0, S_0
else
  | **return** v_1, S_1

Proposition 5. *Alg. 2 returns a best response.*

The correctness of Prop. 5 follows from the following optimal substructure property, which can be easily proved using properties of the Mallows distribution.

Claim 6. *Let $\mathbf{R^{-j}}$ be the strategy profile of all but vendor j. If $S_{st} \subseteq \{c_{j1}, \ldots, c_j\}$ is the revenue maximizing set of size s consisting of items in $\{c_{j1}, \ldots, c_{jt}\}$, and furthermore, $c_{jt} \in S_{st}$ then $S_{st} \setminus \{c_{jt}\}$ is the revenue maximizing set of size $s-1$ consisting of items from $\{c_{j1}, \ldots, c_{j(t-1)}\}$.*

Mallows mixtures The modeling power of the Mallows distribution can be extended by considering mixtures of such models, e.g., reflecting a population with several diverse *types* of consumers.

A *Mallows mixture* is given by (a) d Mallows distributions $D_1(\hat{\pi}_1, \varphi_1), \ldots, D_d(\hat{\pi}_d, \varphi_d)$ and (b) a vector $\mathbf{q} = (q_1, \ldots, q_d)$ ($q_j \in (0,1)$, and $\sum_{t=1}^d q_t = 1$). Each preference σ is sampled i.i.d. from the mixture distribution by first selecting a distribution D_t with probability q_t, and then sampling a ranking using $D_t(\hat{\pi}_t, \varphi_t)$.

Alg. 2 can be modified to handle Mallows mixtures as follows. We first sort the items in $C^j = \{c_{j1}, \ldots, c_{jm_j}\}$ based on their *weighted ranks*, where the weighted rank of $c \in C^j$ is $\overline{\pi}(c) = \sum_{d=1}^d q_t \cdot \hat{\pi}_t(c)$. Similarly, vendor j's revenue w.r.t. catalog $S \subseteq C^j$ is defined using a linear combination of the revenues for each Mallows component:

$$r_j(S, \mathbf{R^{-j}}) = \sum_{t=1}^d r_j^t(S, \mathbf{R^{-j}}) = \sum_{t=1}^d q_t \sum_{c \in S} \frac{\varphi_t^{\ell^j + |S| - 1}}{Z^t_{\ell^j + |S|}},$$

where the normalizing term is $Z^t_m = \sum_{q=1}^m \varphi_t^{q-1}$.

6 Equilibria and Stability

We have seen that (deterministic) best response computation is difficult in some cases, and easy in others. We now turn our attention to the existence pure Nash equilibria and analyze their welfare properties. We first consider games with disjoint vendor sets, then examine the special case where all vendors have identical products from which to choose.

6.1 Disjoint Vendor Sets

In this section we assume all vendor sets C^j are disjoint. We assume familiarity with Nash equilibria, but briefly, a *pure Nash equilibrium (PNE)* in our setting is a vendor strategy profile $\mathbf{R} = (R^1, \ldots, R^k)$ s.t. R^j is a best response to $\mathbf{R^{-j}}$, for each $j \leq k$. A PNE is a *stable* solution in which each vendor offers a catalog that maximizes her revenue given the catalogs of all other vendors. While in any finite normal form game such as ours, Nash equilibria are guaranteed to exist in *mixed strategies* (i.e., where vendors may randomize their choice of catalog), it is not *a priori* clear that our catalog selection games always admits *pure* equilibria.

In the full information case, there are games where all best response paths are cyclic, hence there is no PNE:

Claim 7. *There are instances of full-information catalog selection games which admit no pure Nash equilibrium.*

Proof. A simple counterexample suffices: consider two vendors 1 and 2, with $C^1 = \{a_1, a_2\}, C^2 = \{b_1, b_2\}$. Let $p(a_1) = 2x, p(a_2) = x + \epsilon$, for some $x > 0$, and $0 < \epsilon < x$. Similarly, let $p(b_1) = 2y, p(b_2) = y + \epsilon'$, for $y > 0$ and $0 < \epsilon' < y$. Assume three consumers with preferences:

$$a_2 \succ_1 b_2 \succ_1 a_1 \succ_1 b_1$$
$$b_2 \succ_2 b_1 \succ_2 a_2 \succ_2 a_1$$
$$a_2 \succ_3 a_1 \succ_3 b_1 \succ_3 b_2$$

Vendor 1's best response includes a_2 in R^1 iff 2 includes b_2 in R^2. On the other hand, 2's best response includes b_2 iff 1 does *not* include a_2. This shows the game has no PNE. □

Lack of PNE can occur even when one restricts vendor strategies. For instance, if vendors are limited to catalogs of size 1, one can construct games where no PNE exist.

Special case: single-unit strategies When imposing single unit strategies ($|R^i| = 1$), there may not be a Nash equilibrium. Consider the scenario where there are 2 players, 1 and 2 with sets $C^1 = \{a, a'\}, C^2 = \{b, b'\}$, where all of the prices are 1. Now, consider consumers with the following preferences over $C^1 \cup C^2$ (the valuations can be set so as to induce such rankings):

$$\sigma^{(1)} = a \succ b \succ a' \succ b'$$
$$\sigma^{(2)} = b' \succ a \succ b \succ a'$$
$$\sigma^{(3)} = a' \succ b' \succ a \succ b$$
$$\sigma^{(4)} = b \succ a' \succ b' \succ a$$

It can be verified that there is no stable (pure) strategy profile in this case.

The counterexamples to PNE above rely on precise vendor knowledge of the preference profile. We now anlayze the partial information game, assuming consumer preferences are drawn i.i.d. from a Mallows model (π, φ).

If $\varphi = 0$ (i.e., all consumers have the same preference σ), the game clearly admits a single (type of) PNE: Let $\pi^{-1}(1), \ldots, \pi^{-1}(t)$ be the longest prefix of π whose items belong to a single vendor j; j's dominant strategy is to select only items in this set with maximal price. The revenue of any other vendor is 0, regardless of her strategy.

Now consider the impartial culture model, where $\varphi = 1$. W.l.o.g., we relabel each C^j so they are ordered in non-increasing order of price: $p(c_{j1}) \geq p(c_{j2}) \cdots \geq p(c_{jm_j})$. Using simple best-response dynamics, we will show that a PNE exists. First, by properties of IC, we have:

Observation 8. *Given $\mathbf{R}^{-\mathbf{j}}$, vendor j' (expected) revenue maximizing catalog of size t is $\{c_{j1}, \ldots, c_{jt}\}$.*

This can be shown by noticing that if j selects another catalog R^j of size t, replacing any item $c \in R^j \setminus \{c_{j1}, \ldots, c_{jt}\}$ with an item $c' \in \{c_{j1}, \ldots, c_{jt}\} \setminus R^j$ only (non-strictly) increases j's expected revenue (with no change only if both have equal prices). Obs. 8 implies that specifying the size of a vendor's best response immediately determines the maximal profit achievable by a catalog of this size. While best-response computation need not be tractable to prove the *existence* of a PNE, Obs. 8 implies that we can use best-response dynamics (Alg. 3) to efficiently compute a PNE.

Claim 9. *Let $\mathbf{R}$ and $\mathbf{T}$ be strategy profiles s.t. (a) R^j and T^j are best responses to $\mathbf{R}^{-j}$ and $\mathbf{T}^{-j}$, resp., as in Alg. 3,; and (b) $\sum_{j' \neq j} |R^{j'}| \leq \sum_{j' \neq j} |T^{j'}|$. Then $|R^j| \leq |T^j|$.*

Proof. Consider two profiles $\mathbf{R}, \mathbf{T}$, and let $d = \sum_{j' \neq j}(|T^{j'}| - |R^{j'}|)$. Suppose the claim is false, so $|T^j| = \ell < |R^j|$. By Obs. 8 we have:

$$r_j(T^j, \mathbf{T}^{-\mathbf{j}}) = \frac{\sum_{c \in T^j} p(c)}{\ell + \sum_{j' \neq j} |R^{j'}| + d} = \frac{\sum_{t=1}^{\ell} p(c_{jt})}{\ell + \sum_{j' \neq j} |R^{j'}| + d}$$
$$\geq \frac{\sum_{t=1}^{|R^j|} p(c_{jt})}{|R^j| + \sum_{j' \neq j} |R^{j'}| + d} = r_j(R^j, \mathbf{T}^{-\mathbf{j}})$$

(the inequality follows by best response). This means:

$$\sum_{t=1}^{\ell} p(c_{jt}) \cdot (|R^j| + \sum_{j' \neq j} |R^{j'}|) \geq$$
$$\sum_{t=1}^{|R^j|} p(c_{jt}) \cdot (\ell + \sum_{j' \neq j} |R^{j'}|) + d \cdot (\sum_{t=1}^{|R^j|} p(c_{jt}) - \sum_{t=1}^{\ell} p(c_{jt}))$$
$$> \sum_{t=1}^{|R^j|} p(c_{jt}) \cdot (\ell + \sum_{j' \neq j} |R^{j'}|)$$

where the last inequality follows from $|R^j| > \ell$. This implies $\frac{\sum_{t=1}^{\ell} p(c_{jt})}{\ell + \sum^{j' \neq j} |R^{j'}|} > \frac{\sum_{t=1}^{|R^j|} p(c_{jt})}{|R^j| + \sum^{j' \neq j} |R^{j'}|}$, a contradiction. □

Algorithm 3: Best response dynamics under IC

Initialize: $R^j \leftarrow \emptyset$
while $(R^1, \ldots, R^k)$ *is not a pure Nash equilibrium* **do**
  Let $j \in [k] = \{1, \ldots, k\}$ with a profitable deviation.
  Let $t = \max_{R \subseteq C^j} r_j(R, \mathbf{R^{-j}})$.
  Set $R^j = \arg\max_{R: r_j(R, \mathbf{R^{-j}}) = t}(|R|)$.

Claim 9 implies that the size of best-response catalogs can only increase during execution of best response dynamics, leading to a PNE.

Proposition 10. *Any catalog selection game in which vendors sets are disjoint and consumer preferences are drawn from an impartial culture admits a pure Nash equilibrium.*

Furthermore, the number of iterations of Alg. 3 is at most $O(m)$. Since a vendor requires at most $O(m)$ steps to find a best response, a PNE can be computed in $O(m^2)$ time.

Given that PNE exist under Mallows distributions when $\varphi = 0$ and $\varphi = 1$, we might hope this holds in the intermediate cases where $\varphi \in (0, 1)$. We have yet to resolve this, but conjecture the following, rather strong monotone-consistency property:

Conjecture 11. *Let $\mathcal{G}_1 = (C^1, \ldots, C^k, \mathbf{p}, N, D_1)$ and $\mathcal{G}_2 = (C^1, \ldots, C^k, \mathbf{p}, N, D_2)$ be two catalog selection games that differ only in the dispersion of their underlying Mallows distributions $D_1(\pi, \varphi_1), D_2(\pi, \varphi_2)$, with $\varphi_1 > \varphi_2$. Then if G_1 admits a PNE, so does G_2.*

If the conjecture is true, it implies the existence of PNE for every value of φ.

While games under IC admit PNE, some of these equilibria may be extremely inefficient due to asymmetries in vendor item sets. As an efficiency metric, we use *vendor social welfare*, $sw_{\mathcal{G}}(\mathbf{R})$ (i.e., expected total vendor revenue given the preference prior under profile $\mathbf{R}$), and both the *price of anarchy* and the *price of stability*:

Definition 12. *Let of $\mathcal{H}$ be the set of games with k vendors with disjoint item sets, m total items, and IC priors. The* price of anarchy (PoA)*, and the* price of stability (PoS) *are:*

$$PoA = \max_{\mathcal{G} \in \mathcal{H}} \frac{\max_{\mathbf{R} \in 2^{C^1} \times \cdots 2^{C^k}} sw_{\mathcal{G}}(R)}{\min_{\substack{\mathbf{R}: \\ \mathbf{R} \text{ is a PNE}}} sw_{\mathcal{G}}(R)}$$
$$PoS = \max_{\mathcal{G} \in \mathcal{H}} \frac{\max_{\mathbf{R} \in 2^{C^1} \times \cdots 2^{C^k}} sw_{\mathcal{G}}(R)}{\max_{\substack{\mathbf{R}: \\ \mathbf{R} \text{ is a PNE}}} sw_{\mathcal{G}}(R)}$$

PoA (PoS) is the worst-case ratio of optimal, non-strategic social welfare realizable by *any* strategy profile to the worst (best) social welfare in *some PNE*. Both PoA and PoS can grow linearly with the number of items:

Claim 13. *There are catalog selection games with partial information in which the PoA and PoS are both* $\Theta(m)$.

Proof. Consider a game with two vendors, with $C^1 = \{c_{11}\}$ and $C^2 = \{c_{21}, \ldots, c_{2T}\}$, for some T. Let $p(c_{11}) = 1$, and $p(c_{21} = \ldots = p_{2T}) = \epsilon = O(1/m)$. Assume consumer preferences are IC. Clearly, the only PNE has both vendors select all items. Since each consumer selects item c_{11} with probability $1/(T+1)$ and some item worth ϵ with probability $T/(T+1)$, the claim follows. □

While this PNE is highly inefficient from the vendors' perspective, it is very efficient from the consumers' side, since it allows them to choose more desirable items.

6.2 Vendors with identical sets

As shown above, one reason for the inefficiency of some equilibria stems from asymmetry in the item sets. It is thus interesting to consider the other extreme case, where $C^1 = \ldots = C^k$. In both the full and partial information settings it is easy to see that PNE always exist:

Observation 14. *Any instance of the catalog selection game with identical vendor item sets admits a PNE.*

This can be verified by noticing that if each vendor offers the entire set C, no vendor benefits by deviating. Moreover, as discussed above, if an item c is selected by some vendor, all vendor best responses must include c.

In the full information case, there are instances in which the only (hence, best) PNE is highly inefficient:

Claim 15. *There are full information games with common item sets in which the PoS is* $\Omega(2^m)$.

Proof. Consider a game with two vendors, $C = \{c_1, \ldots, c_m\}$, and prices: $p(c_1) = 1$; $p(c_i) = \frac{c_{i-1}}{2} + \epsilon$ for $i \geq 2$, where $\epsilon = e^{-m}/2$. Assume a single consumer with preference $c_m \succ \cdots \succ c_1$. In any PNE, all vendors select item c_m. Hence, the revenue in the best PNE is $p(c_m) = 2^{-(m-1)} + O(e^{-m}) = \Theta(2^{-m})$, in contrast to the optimal total revenue of 1. □

We also consider the partial information case under IC. Given a strategy profile $\mathbf{R}$, and letting $A = \bigcup_{i=1}^{k} R^i$, the revenue of vendor j is $r_j(R^j, \mathbf{R^{-j}}) = \frac{1}{k\cdot|A|}\sum_{c\in R^j\cap T} p(c) + \frac{1}{|A|}\sum_{c\in R^j\setminus T} p(c)$.

Theorem 16. *Given identical item sets, if preferences are drawn from IC, then PoA is* $\Theta(m)$.

Proof. Consider a game with 2 vendors, $C = \{c_1, \ldots, c_m\}$, and prices: $p(c_1) = 1, p(c_i) = \epsilon, i \geq 1$. Consider a PNE where all vendors select C. The revenue in is $\frac{\sum_{i=1}^{m} p(c_i)}{m} = \frac{1}{m} + \frac{(m-1)\epsilon}{m}$. The optimal total revenue is 1. □

If the number of vendors is assumed to be constant, then PoS is a logarithmic factor smaller than PoA.

Theorem 17. *Given identical item sets, if preferences are drawn from IC, then PoS is* $\Theta(\frac{m\cdot k}{\log m})$.

Let $r_j(X, \mathbf{Y^{k-1}})$ be vendor j's utility when selecting set X in response to competitor profile Y. Furthermore, let $P_i = \{c_1, \ldots, c_i\}$. The following lemma shows that Algorithm 4 always returns a PNE:

Lemma 18. *If the algorithm halts at step* $i < m$*, then* $(P_i, \ldots, P_i)$ *is a Nash equilibrium.*

Algorithm 4: Finding a Nash equilibrium

Input: k vendors, items $C = \{c_1, \ldots, c_m\}$, price vector $\mathbf{p}$ such that $p(c_1) \geq \cdots \geq p(c_m)$

for $i \leftarrow 2$ *to* m **do**

  if $r_1(P_{i-1}, \mathbf{P^{k-1}_{i-1}}) \geq r_1(P_{i-1} \cup \{c_i\}, \mathbf{P^{k-1}_{i-1}})$ **then**

    return P_{i-1}

return C

Proof. As the items are ordered in a non-increasing order of price, it suffices to show that no (arbitrary, due to symmetry) vendor would deviate by selecting a prefix P_j, for $j > i$. We show inductively that if a vendor improves by deviating to such P_j, then she can do so by deviating to P_{i+1} as well. W.l.o.g., assume the first vendor deviates. First, we show that if a vendor improves her revenue by selecting P_j then she can improve it by deviating to P_{j-1}. Suppose by way of contradiction that $r_1(P_j, \mathbf{P^{k-1}_i}) > r_1(P_i, \mathbf{P^{k-1}_i})$, for $j > i$, but $r_1(P_{j-1}, \mathbf{P^{k-1}_i}) \leq r_1(P_i, \mathbf{P^{k-1}_i})$. Then by definition

$$\frac{\sum_{t=1}^{i} p(c_t)}{k \cdot i} \geq \frac{\sum_{t=1}^{i} p(c_t)}{k \cdot (j-1)} + \frac{\sum_{t=i+1}^{j-1} p(c_t)}{j-1},$$

which implies $\sum_{t=i+1}^{j-1} p(c_t) \leq \frac{j-i-1}{i \cdot k} \sum_{t=1}^{i} p(c_t)$. Then:

$$\begin{aligned}
r_1(P_j, \mathbf{P^{k-1}_i}) &= \frac{\sum_{t=1}^{i} p(c_t)}{k \cdot j} + \frac{\sum_{t=i+1}^{j-1} p(c_t)}{j} + \frac{p(c_j)}{j} \\
&\leq \frac{\sum_{t=1}^{i} p(c_t)}{k \cdot j} + \frac{(j-i-1)\sum_{t=1}^{i} p(c_t)}{j \cdot k \cdot i} + \frac{\sum_{t=i+1}^{j-1} p(c_t)}{j \cdot (j-i-1)} \\
&\leq \frac{\sum_{t=1}^{i} p(c_t)}{k \cdot j} + \frac{(j-i-1)\sum_{t=1}^{i} p(c_t)}{j \cdot k \cdot i} + \frac{\sum_{t=1}^{i} p(c_t)}{j \cdot k \cdot i} \\
&= \frac{\sum_{t=1}^{i} p(c_t)}{k \cdot i} = r_1(P_i, \mathbf{P^{k-1}_i})
\end{aligned}$$

where the first inequality follows from the bound above and an averaging argument on $p(c_j)$. This is a contradiction. Hence, deviating to P_{j-1} also improves vendor revenue. Repeating this process until $i+1$ contradicts the stopping condition of the for-loop of the algorithm. □

Next, we bound the rate of decrease in prices to construct a lower bound on expected social welfare.

Lemma 19. *Suppose Alg. 4 returns set* $P_i = \{c_1, \ldots, c_i\}$. *Then* $p(c_j) \geq \frac{1}{k \cdot (j-1)} + \Theta(\frac{1}{k^2})$, *for* $2 \leq j \leq i$.

Proof. The algorithm stops when $r_1(P_i, \mathbf{P^{k-1}_{i-1}}) \leq r_1(P_{i-1}, \mathbf{P^{k-1}_{i-1}})$. Using the definitions of $r_1(P_i, \mathbf{P^{k-1}_{i-1}})$ and $r_1(P_{i-1}, \mathbf{P^{k-1}_{i-1}})$, and rearranging the terms, we get that for every $1 < j \leq i$,

$$\frac{\sum_{t=1}^{i} p(c_t)}{(k \cdot (j-1)} < \frac{\sum_{t=1}^{j-1} p(c_t)}{k \cdot j} + \frac{p(c_j)}{j}$$

which implies the recursive inequality: $p(c_j) > \frac{\sum_{t=1}^{j-1} p(c_t)}{k \cdot (j-1)}$. The statement of the lemma can be then shown to be the solution of this inequality, using induction. □

Proof of Thm. 17. The worst case execution of Alg. 4 occurs when it reaches the last item. By Lemma 19, expected welfare is bounded below by $\frac{1}{m}(1 + \sum_{i=2}^{m} \frac{1}{k \cdot (i-1)}) = \Omega(\frac{\ln m}{m \cdot k})$. The fact that $p(c_1) = 1$ implies the upper bound on PoS. We can construct a matching worst-case price vector using the bound on the $p(c_i)$'s given in Lemma 19. □

7 Conclusions

We have presented a model of competition among vendors who offer slates or catalogs of products to their consumers using rank-based models of preferences that have connections to models in computational social choice and algorithmic game theory. We studied both best response computation (and equlibrium finding in some cases) and various equilibrium properties under two different informational assumptions w.r.t. conumser preferences.

There are a number of interesting directions remaining to be explored. The possibility of approximating best responses in the full information setting remains open. This problem does not appear to have any of the usual "nice" properties often used for devising efficient optimization algorithms (e.g., symmetry, monotonicity, submodularity). The study of our model in cases where the strategies are required to satisfy certain combinatorial constraints (e.g., matroids or knapsack constraints), reflecting limits on individual catalogs, would be of interest. Under some such restrictions, our worst case PoA and PoS ratios might be improved. Connections to other models in game theory also bears exploration. For instance, allowing for endogenous prices gives a framework where vendors must post prices in a multi-vendor platform, offering a competitive extension of the well-studied area of profit-maximizing, envy-free mechanisms (see e.g., [6]).

References

[1] D. Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.

[2] J. R. Chamberlin and P. N. Courant. Representative deliberations and representative decisions: Proportional representation and the borda rule. *The American Political Science Review*, 77(3):pp. 718–733, 1983.

[3] D. Cornaz, L. Galand, and O. Spanjaard. Bounded single-peaked width and proportional representation. In *ECAI*, pages 270–275, 2012.

[4] U. Feige. A threshold of ln n for approximating set cover. *J. ACM*, 45(4):634–652, 1998.

[5] A. V. Goldberg, J. D. Hartline, A. R. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.

[6] V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, pages 1164–1173, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.

[7] D. Honhon, S. Jonnalagedda, and X. A. Pan. Optimal algorithms for assortment selection under ranking-based consumer choice models. *Manufacturing & Service Operations Management*, 14(2):279–289, 2012.

[8] J. Kleinberg, C. Papadimitriou, and P. Raghavan. Segmentation problems. *Journal of the ACM*, 51:263–280, 2004.

[9] X. Li, S. Nukala, and S. Mohebbi. Game theory methodology for optimizing retailers' pricing and shelf-space allocation decisions on competing substitutable products. *The International Journal of Advanced Manufacturing Technology*, 68(1-4):375–389, 2013.

[10] T. Lu and C. Boutilier. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the Twenty-Second international joint conference on Artificial Intelligence-Volume Volume One*, pages 280–286. AAAI Press, 2011.

[11] C. L. Mallows. Non-null ranking models. *Biometrika*, 44:114–130, 1957.

[12] J. I. Marden. *Analyzing and Modeling Rank Data*. Chapman and Hall, London, 1995.

[13] V. Martínez-de Albéniz and G. Roels. Competing for shelf space. *Production and Operations Management*, 20(1):32–46, 2011.

[14] B. L. Monroe. Fully proportional representation. *The American Political Science Review*, 89(4):pp. 925–940, 1995.

[15] C. Schön. On the product line selection problem under attraction choice models of consumer behavior. *European Journal of Operational Research*, 206(1):260–264, October 2010.

[16] P. Skowron, P. Faliszewski, and A. Slinko. Fully proportional representation as resource allocation: Approximability results. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pages 353–359. AAAI Press, 2013.

Joel Oren
University of Toronto
Toronto, Canada
Email: `oren@cs.toronto.edu`

Nina Narodytska
University of Toronto
Toronto, Canada
Email: `ninan@cs.toronto.edu`

Craig Boutilier
University of Toronto
Toronto, Canada
Email: `cebly@cs.toronto.edu`