

# A Game-theoretic Analysis of Catalog Optimization

Joel Oren and Nina Narodytska and Craig Boutilier

Department of Computer Science, University of Toronto

{oren, ninan, cebly}@cs.toronto.edu

## 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

## 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 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 equilibria exist 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 in the total number of products, and provide additional analysis of special cases in which all vendors have identical product sets.

## Preliminaries

We consider a game with  $k$  strategic vendors, each vendor  $j$  having access to a set  $C^j = \{c_{j1}, \dots, c_{jm_j}\}$  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  $C = \bigcup_{j=1, \dots, k} C^j$ , 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  $\mathbf{R} = (R^1, \dots, R^k)$ .

We assume a set of *unit-demand* consumers  $N = \{1, \dots, n\}$ . Each consumer  $i$  has a *strict preference ordering*  $\pi_i \in \mathcal{L}_m$ , representing her preferences over items  $C$ . If  $\pi_i(c) < \pi_i(c')$ , then  $i$  prefers  $c$  to  $c'$  (given their fixed prices); i.e.  $i$  ranks  $c$  above  $c'$ . While we focus on ordinal preferences,  $\pi_i$  can be a ranking induced by  $i$ 's intrinsic valuation for items and the prices. Let  $\boldsymbol{\pi} = (\pi_1, \dots, \pi_n)$  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.<sup>1</sup> Given a (non-empty) set of offered items  $A = \bigcup_{j=1}^k R^j$ , each  $i \in N$  her their most preferred item in  $A$ , i.e.,  $\text{top}_i(A) = \arg \min_{c \in A} \pi_i(c)$ . Let  $\text{top}_i(\mathbf{R}) = \text{top}_i(\bigcup_{j=1}^k R^j)$ . This determines a game  $\mathcal{G} = (C^1, \dots, C^k, \mathbf{p}, N, \boldsymbol{\pi})$ , in which vendor  $j$ 's strategies are catalogs  $R_j \subseteq C^j$ , and her revenue (payoff) is determined by the *strategy profile*  $\mathbf{R}$ . If the  $C^j$  are not disjoint and consumer  $i$  selects a  $c$  that occurs in several vendor catalogs, 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:

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

Vendor  $j$ 's *best response* to  $\mathbf{R}^{-j}$  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  $\boldsymbol{\pi}$ . In the *partial information* setting, vendors instead have only a common prior belief, or a *distribution* over profiles. We assume consumer preferences are drawn i.i.d. from a *Mallows  $\varphi$ -distribution* (or a mixture thereof), a probabilistic model of rankings widely used in statistics, machine learning, econometrics and social choice (Mallows 1957; Marden 1995). The Mallows model is specified by two parameters, a *reference ranking*  $\hat{\pi} \in \mathcal{L}_m$  and a *dispersion parameter*  $\varphi$  (controlling variance). The probability of a ranking  $\pi$  in this model is  $\Pr(\pi) = \varphi^{\tau(\pi, \hat{\pi})} / T_m$ , where  $\tau$  is the Kendall-tau (or swap) distance between two rankings, and  $T_m = \prod_{t=1}^m Z_t$  is a normalization constant, 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.

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

**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., Schön (2010)), where the aim is to select a revenue-maximizing catalog assuming some consumer preference model. More relevant is work addressing the *strategic aspects* of this problem in the face of competition, assessing both computation and existence of equilibria. Martínez-de Albéniz and Roels (2011) study the efficiency of equilibria of vendors competing for *limited* shelf space (under both exogenous and endogenous prices). Our approach differs w.r.t. the consumer preference model and shelf-space constraints. Li et al. (2013) similar models. Hohnon et al. (2012) address (non-competitive) variants of the problem using a rank-based preference model similar to ours.

Our model can also be thought of an extension *multi-winner social choice (MWSC)*. In MWSC, the goal is to select a “slate” of items given a set of agent preferences, and has application to legislature/committee selection (Chamberlain and Courant 1983; Monroe 1995), facility location, and group (e.g., consumer) decision-making (Kleinberg, Papadimitriou, and Raghavan 2004; Lu and Boutilier 2011; Skowron, Faliszewski, and Slinko 2013). 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 (Goldberg et al. 2006; Guruswami et al. 2005).

## 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.<sup>2</sup> 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 Max-SNP hard (Feige 1998). Take as input  $m$  DNF clauses  $(\phi_1, \dots, \phi_m)$  over  $n$  variables  $(x_1, \dots, x_n)$ , where each variable is contained in exactly five clauses. For 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  $\{a_i, b_i\}_{i \leq n}$  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  $x_1, \dots, x_n$ , i.e., that exactly

<sup>2</sup>This is not the case if preferences are truncated.

one of  $\{a_i, b_i\}$  is *True*. For  $i \leq n$  create partial rankings:

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

(The competitors' items are always ranked below the items in  $C^j$ ). Including  $t$  can never hurt, so we assume it is always chosen. If both or neither of  $a_i$  and  $b_i$  are selected, payoff is 3. If only one is selected, payoff is 3.5.

**Set 2: Clause rankings.** For each clause of the form  $\phi_j = \ell_{j1} \vee \ell_{j2} \vee \ell_{j3}$ , create a ranking:  $\pi_j : f(\ell_{j1}) > f(\ell_{j2}) > f(\ell_{j3})$ , where  $f(\ell_{ji}) = a_{ji}$  if  $\ell_{ji}$  corresponds to variable  $x_{ji}$  in its non-negated form, and  $b_{ji}$  otherwise. If  $\phi_j$  is satisfied, at least one item corresponding to  $\ell_{j1}, \ell_{j2}$ , and  $\ell_{j3}$  is selected, which gives an extra payoff of 1 from ranking  $\pi_j$ .

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  $\phi$  is only  $(1 - \epsilon)$ -satisfiable, the maximum value we can obtain is  $(1 - \epsilon)m + 3.5n = (1 - \epsilon)m + 21m/10$ , hence it is NP-hard to distinguish between cases with a profit of  $m + 3.5n$  from cases with only a  $(1 - \delta)$  fraction of that  $m + 3.5n$  profit, where  $\delta = \epsilon/(1 + 21/10)$ .  $\square$

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  $\beta$ -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. In the following section, we consider a restriction in which agent preferences are *single-peaked*, which we address in the next section.

## Single-Peaked, Truncated Preferences

We now consider an important restriction on the preferences of consumers. Suppose consumers are *single-peaked*, which we define below. The class of single-peaked preferences has been deeply studied in social choice (Black 1948), and is well-suited for modeling preferences in a variety of domains (including those based on distances). The single-peaked assumption has a variety of attractive computational and incentive properties as well (Gaertner 2002).

In addition to single-peakedness, we impose an additional restriction on preferences that interacts with the items being offered. Consider a setting in which any individual consumer  $i$ 's preference tends to be such that the items of a number of different vendors are near its peak. Even further, suppose that no matter what  $i$ 's most preferred item is from the set  $\mathbf{R}$  of those actually offered by all vendors, there are items offered by competing vendors that are reasonably highly ranked as well. This would be the case, for example, in reasonably competitive environments in which, for any product offered by some vendor  $j$ , there is a competitor  $j'$  with a similar product at a similar price point. This would also be the case when most vendors have small sets of items; no

one vendor can “dominate” too many of the top positions of  $i$ 's preference ranking. In such cases, consumer preferences are effectively *truncated*; i.e., from the perspective of any vendor  $j$ , each consumer  $i$ 's preference will have an offered item from at least one *competitor*  $j'$  that is highly ranked. This means any of  $j$ 's items falling below this competitive item in  $i$ 's ranking will never be selected by  $i$ , even if offered.

We now formally define the notion of *truncated, single-peaked preference profiles* (combining both notions):

**Definition 2.** For vendor  $j$ , and strategy profile  $\mathbf{R}^{-j}$  for other vendors, preference profile  $(\pi_1, \dots, \pi_n)$  is *single-peaked, L-truncated* if:

1. For every  $i \in N$ , letting  $t_i = \arg \min_{c \in \cup_{j' \neq j} R^{j'}} \pi_i(c)$ , we have that  $t_i \leq L + 1$ . For convenience, define  $S_i^j = \{c \in C^j : \pi_i(c) < t_i\}$  to be those items in  $C^j$  that are “relevant” to  $i$ .
2. There is an ordering  $\sigma \in \mathcal{L}_m$  of the items in  $C$  (i.e., an axis) s.t. for each  $i \in N$ : there is some  $c \in C$  (the peak) s.t. for any  $c', c'' \in S_i^j \setminus \{c\}$ , if  $\sigma(c) > \sigma(c') > \pi(c'')$  or  $\sigma(c) < \sigma(c') < \sigma(c'')$ , then  $\pi_i(c') < \pi_i(c'')$ .

The first condition ensures that there are at most  $L$  relevant items per consumer in  $C^j$ . The second is the usual single-peaked condition, except that we only impose the restriction on the *preference prefixes* of  $j$ -relevant items (not the whole set  $C$  or  $C^j$ ).

Assume that items in  $C^j$  are labeled according to their order along the axis  $\sigma$  (i.e., for  $c_{jt}, c_{j\ell} \in C^j$ ,  $t < \ell$  iff  $\sigma(c_{jt}) < \sigma(c_{j\ell})$ ). Recall that we can assume w.l.o.g. that no items in  $C^j$  are also offered by a competitor. The following theorem shows that optimizing a vendor's best-response is tractable under these two conditions.

**Theorem 3.** Let  $\mathbf{R}^{-j}$  be a strategy profile of all vendors except  $j$ . If a preference profile is single-peaked, L-truncated, then  $j$ 's best response to  $\mathbf{R}^{-j}$  can be computed in  $O(2^L nm^2)$  time.

The following claim is instrumental in efficiently finding a best response for a given vendor.

**Claim 4.** If  $i \in N$  has a single-peaked L-truncated preference  $\pi_i$  over  $S_i^j \subseteq C^j$ , then  $\max_{c, c' \in C^j} |\sigma(c) - \sigma(c')| < L$ .

*Proof.* Suppose consumer  $i$  has such a preference, but there are two  $c', c'' \in S_i^j$  s.t.  $|\sigma(c') - \sigma(c'')| > L$ . By the pigeon-hole principle there must be a  $c^* \in C \setminus S_i^j$  that lies between  $c'$  and  $c''$  on the axis  $\sigma$ . This contradicts the fact that  $i$ 's preference is single-peaked.  $\square$

Intuitively, Claim 4 implies that the decision to include (or not) an item  $c$  cannot affect consumers with truncated preferences over sets that contain items distance greater than  $L$  from  $c$  w.r.t.  $\sigma$ .

We now describe a dynamic programming algorithm, see Alg. 1, for optimizing a vendor's best-response. Let  $L \leq t \leq m$ . Suppose that for each subset  $S'$  of items in  $\{\sigma^{-1}(t - L), \dots, \sigma^{-1}(t - 1)\}$  we have computed the optimal slate that includes *all* items in  $S'$  and some subset

of  $\{\sigma^{-1}(1), \dots, \sigma^{-1}(t-L-1)\}$ . To compute the optimal value for some  $S \subseteq \{\sigma^{-1}(t-L+1), \dots, \sigma^{-1}(t)\}$  we only need to consider at most two subsets:  $S \setminus \{\sigma^{-1}(t)\}$ , and  $S \setminus \{\sigma^{-1}(t)\} \cup \{\sigma^{-1}(t-L)\}$ .

---

**Algorithm 1:** The best response finding algorithm for single-peaked,  $L$ -truncated preferences.

---

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

```

1 Notation: For a binary vector  $v \in \{0, 1\}^L$ , we let  $B(v)$ 
   denote the decimal representation of  $v$ :  $B(v) = \sum_{t=1}^k 2^{t-1}$ .
2 For a subset  $S \subseteq \{c_{j(t-L+1)}, \dots, c_{jt}\}$  we let  $\mathbb{1}_{S,t}$  denote the
   length- $L$  characteristic vector of  $S$  w.r.t.
    $\{c_{j(t-L+1)}, \dots, c_{jt}\}$ . That is,  $\mathbb{1}_{S,t}(d) = 1$  iff  $c_{j(t-L+d)} \in S$ 
   and 0 otherwise, for  $d = 1, \dots, L$ .
3 Let  $M$  be an  $m_j$  by  $2^L$  table.
4 For  $t \in \{1, \dots, m_j\}$  and  $S \subseteq \{c_{j(t-L+1)}, \dots, c_{jt}\}$ ,
    $M[t, B(\mathbb{1}_{S,t})]$  contains the optimal solution for the problem
   of optimizing the state using items from  $\{c_{j1}, \dots, c_{jt}\}$ , such
   that  $M[t, B(\mathbb{1}_{S,t})] \setminus \{c_{j1}, \dots, c_{j(t-L)}\} = S$ .
5 for  $t \rightarrow 1$  to  $L$  do
6   foreach  $S \subseteq \{c_{j1}, \dots, c_{jt}\}$  do
7      $M[t, B(\mathbb{1}_{S,t})] \leftarrow S$ 
8 for  $t \leftarrow L+1$  to  $m_j$  do
9   foreach  $S \subseteq \{c_{j(t-L+1)}, \dots, c_{jt}\}$  do
10     $S_1 \leftarrow M[t-1, B(\mathbb{1}_{S \setminus \{c_{jt}\}, t-1})]$ 
11     $S_2 \leftarrow M[i-1, B(\mathbb{1}_{S \setminus \{c_{jt}\} \cup \{c_{j(t-L)}\}, t-1})]$ 
12    if  $r_j(S_1 \cup S, \mathbf{R}^{-j}) \geq r_j(S_2 \cup S, \mathbf{R}^{-j})$  then
13       $M[t, B(\mathbb{1}_{S,t})] \leftarrow S_1 \cup S$ 
14    else
15       $M[t, B(\mathbb{1}_{S,t})] \leftarrow S_2 \cup S$ 
16 return  $\arg \max_{S \subseteq \{c_{j(m_j-L+1)}, \dots, c_{jm_j}\}} M[m, B(\mathbb{1}_{S,m_j})]$ 

```

---

The correctness of Alg. 1 follows immediately from Claim 4: the decision to include item  $c_{jt}$  doesn't affect the revenue derived from consumers whose peaks are further than  $L$  from  $c_{jt}$ , implying the required optimal substructure property. The running time of the algorithm is  $O(2^L n m^2)$ , i.e., polynomial in  $n$  and  $m$  if  $L = O(\log m)$ .

## 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. 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  $C^1, \dots, C^k$  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  $p(c_{j1}) \geq \dots \geq p(c_{jm_j})$ , and define the length  $t$  prefix of this item vector as  $T_t = \{c_{j1}, \dots, c_{jt}\}$  and  $T_0 = \emptyset$ .

For any catalog  $R^j \subseteq C^j$ ,  $j$ '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., greedily add items in decreasing order of price as long as the expected revenue is increased by these additions. Adding *any* other items cannot contribute to  $j$ 's expected profit. The optimality of this algorithm can be proven using an elementary exchange argument.

### Mallows models

We now address a broader class of distributions, preferences drawn i.i.d. from a Mallows model (note that IC is a special case). Vendor  $j$ 's best response can be computed by dynamic programming in this case, see Alg. 2. Assume a competitor strategy profile  $\mathbf{R}^{-j}$ , and beliefs given by Mallows model  $(\hat{\pi}, \varphi)$ . For convenience, 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  $s \leq m_j$  and index  $t = s, \dots, m_j$ , Alg. 2 recursively computes the optimal catalog of size  $s$  consisting of items from subset  $\{c_{j1}, \dots, c_{jt}\}$ . 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  $C^j \setminus R^j$ . We use the recursive nature of the Mallows model to compute  $j$ 's revenue. For  $c_{jt} \in C^j$ , and a set of previously selected items  $S \subseteq \{c_{j1}, \dots, c_{j(t-1)}\}$ , 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  $Z_{\ell^j+s}$  is the normalizing term. Thus, if the expected revenue of selecting  $S$  is  $r_j(S, \mathbf{R}^{-j})$ , that of adding  $c_{jt}$  to  $S$  is:  $r_j(S \cup \{c_{jt}\}, \mathbf{R}^{-j}) = (r_j(S, \mathbf{R}^{-j}) \cdot Z_{\ell^j+s-1} + \varphi^{\hat{\pi}_{S \cup \{c_{jt}\}}-1} \cdot p(c_{jt})) / Z_{\ell^j+s}$ .

**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 a competitor strategy profile. If  $S_{st} \subseteq \{c_{j1}, \dots, c_{jt}\}$  is vendor  $j$ 's revenue maximizing set of size  $s$  consisting of items in  $\{c_{j1}, \dots, c_{jt}\}$  and  $c_{jt} \in S_{st}$ , then  $S_{st} \setminus \{c_{jt}\}$  is the revenue maximizing set of size  $s-1$  consisting of items from  $\{c_{j1}, \dots, c_{j(t-1)}\}$ .

### 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:  $d$  Mallows distributions  $D_1(\hat{\pi}_1, \varphi_1), \dots, D_d(\hat{\pi}_d, \varphi_d)$ ; and a vector  $\mathbf{q} = (q_1, \dots, q_d)$  ( $q_j \in (0, 1)$ , and  $\sum_{t=1}^d q_t = 1$ ). Each preference  $\pi$  is sampled i.i.d. from the mixture distribution by first selecting

---

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

---

```

1 Assume  $\ell^j = \sum_{j' \neq j} |R^{j'}|$ , and for a non-negative integer  $q$ ,
   let  $\tau_q = Z_{\ell^j+q} = \sum_{d=0}^{\ell^j+q-1} \varphi^d$ .
2 for  $s \leftarrow 1$  to  $m_j$  do
3   | Let  $v_s, S_s \leftarrow \text{OptimizeSlate}(s, s)$ 
4   | Return  $S_s$  with maximal value  $v_s$ .
5 OptimizeSlate ( $s, t$ )
   Input: A Mallows distribution with parameters  $(\hat{\pi}, \varphi)$ .
   Output: Optimal slate  $R^j \subseteq \{c_{j1}, \dots, c_{jt}\}$ , s.t.  $|R^j| = s$ .
6 if  $s = 0$  then
7   | return  $0, \emptyset$ 
8 if  $s = t$  then
9   |  $S \leftarrow \{c_{j1}, \dots, c_{jt}\}$ 
10  |  $V \leftarrow \sum_{d=1}^s p(c_{jd}) \cdot \varphi^{\hat{\pi}_S(c_{jd})-1} / \tau_s$ 
11  | return  $S, V$ 
12  $v_0, S_0 \leftarrow \text{OptimizeSlate}(s, t-1)$ 
13  $v_1, S_1 \leftarrow \text{OptimizeSlate}(s-1, t-1)$ 
14  $S_1 \leftarrow S_1 \cup \{c_{jt}\}$ 
15  $v_1 \leftarrow (v_1 \cdot Z_{s-1} + p(c_{jt}) \cdot \varphi^{\hat{\pi}_{S_1}(c_{jt})-1}) / \tau_s$ 
16 if  $v_0 \geq v_1$  then
17   | return  $v_0, S_0$ 
18 else
19   | return  $v_1, S_1$ 

```

---

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  $C^j = \{c_{j1}, \dots, c_{jm_j}\}$  based on their *weighted ranks*, where the weighted rank of  $c \in C^j$  is  $\bar{\pi}(c) = \sum_{d=1}^d q_t \cdot \hat{\pi}_t(c)$ . 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:

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

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

## 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.

### 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  $\mathbf{R} = (R^1, \dots, R^k)$  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:

$$\begin{aligned} a_2 \succ_1 b_2 \succ_1 a_1 \succ_1 b_1, \quad b_2 \succ_2 b_1 \succ_2 a_2 \succ_2 a_1, \\ a_2 \succ_3 a_1 \succ_3 b_1 \succ_3 b_2 \end{aligned}$$

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.  $\square$

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.

The counterexamples to PNE above rely on precise vendor knowledge of the preference profile. We now analyze the partial information game, assuming consumer preferences are drawn i.i.d. from a Mallows model  $(\pi, \varphi)$ . If  $\varphi = 0$  (i.e., all consumers have the same preference  $\pi$ ), the game clearly admits a single (type of) PNE: Let  $\pi^{-1}(1), \dots, \pi^{-1}(t)$  be the longest prefix of  $\pi$  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 impartial culture, where  $\varphi = 1$ . W.l.o.g., relabel each  $C^j$  so they are ordered in non-increasing order of price:  $p(c_{j1}) \geq p(c_{j2}) \dots \geq p(c_{jm_j})$ . Using simple best-response dynamics, we show that a PNE exists. First, by properties of IC, we have:

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

This can be shown by noticing that if  $j$  selects another catalog  $R^j$  of size  $t$ , replacing any item  $c \in R^j \setminus \{c_{j1}, \dots, c_{jt}\}$  with an item  $c' \in \{c_{j1}, \dots, c_{jt}\} \setminus R^j$  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—see 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:

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

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

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

where the last inequality follows since  $|R^j| > \ell$ . 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.  $\square$

---

### Algorithm 3: Best response dynamics under IC

---

- 1 Initialize:  $R^j \leftarrow \emptyset$
  - 2 **while**  $(R^1, \dots, R^k)$  is not a pure Nash equilibrium **do**
  - 3     Let  $j \in [k] = \{1, \dots, k\}$  with a profitable deviation.
  - 4     Let  $t = \max_{R \subseteq C^j} r_j(R, \mathbf{R}^{-j})$ .
  - 5     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 for  $\varphi \in (0, 1)$ . We have yet to resolve this, but conjecture the following, rather strong monotone-consistency property:

**Conjecture 11.** *Let  $\mathcal{G}_1 = (C^1, \dots, C^k, \mathbf{p}, N, D_1)$  and  $\mathcal{G}_2 = (C^1, \dots, C^k, \mathbf{p}, N, D_2)$  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  $\mathcal{G}_1$  admits a PNE, so does  $\mathcal{G}_2$ .*

If true, Conj. 11 implies the existence of PNE for all  $\varphi$ .

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:

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

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  $\Omega(m)$ .*

*Proof.* Consider a game with two vendors, with  $C^1 = \{c_{11}\}$  and  $C^2 = \{c_{21}, \dots, c_{2T}\}$ , for some  $T$ . Let  $p(c_{11}) = 1$ , and  $p(c_{21} = \dots = p_{2T}) = \epsilon = O(1/m)$ . 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  $\epsilon$  with probability  $T/(T+1)$ , the claim follows.  $\square$

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.

### Vendors with identical sets

The inefficiency of some equilibria above stems in part from asymmetry in the item sets. It is thus interesting to consider the other extreme case, where  $C^1 = \dots = C^k$ . 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,  $C = \{c_1, \dots, c_m\}$ , 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  $c_m \succ \dots \succ c_1$ . 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.  $\square$

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,  $C = \{c_1, \dots, c_m\}$ , 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.  $\square$

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  $\Theta(m \cdot k / \log m)$ .*

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  $P_i = \{c_1, \dots, c_i\}$ . Alg. 4 is a simple procedure for comput-

---

**Algorithm 4:** Finding a Nash equilibrium

---

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

```

1 for  $i \leftarrow 2$  to  $m$  do
2   if  $r_1(P_{i-1}, \mathbf{P}_{i-1}^{k-1}) \geq r_1(P_{i-1} \cup \{c_i\}, \mathbf{P}_{i-1}^{k-1})$  then
3     return  $P_{i-1}$ 
4 return  $C$ 

```

---

ing a PNE, as the following lemma shows:

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

*Proof.* As the items are ordered in a non-increasing order of price, it suffices to show that no (arbitrary, by 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 a  $P_j$ , then she can do so by deviating to  $P_{i+1}$  too. Assume w.l.o.g. 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  $r_1(P_j, \mathbf{P}_i^{k-1}) > r_1(P_i, \mathbf{P}_i^{k-1})$ , for  $j > i$ , but  $r_1(P_{j-1}, \mathbf{P}_i^{k-1}) \leq r_1(P_i, \mathbf{P}_i^{k-1})$ . 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  $\sum_{t=i+1}^{j-1} p(c_t) \leq \frac{j-i-1}{i-k} \sum_{t=1}^i p(c_t)$ . Then:

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

where the first inequality follows from the bound above and an averaging argument on  $p(c_j)$ , 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.  $\square$

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  $P_i = \{c_1, \dots, c_i\}$ . Then  $p(c_j) \geq \frac{1}{k \cdot (j-1)} + \Theta(\frac{1}{k^2})$ , for  $2 \leq j \leq i$ .*

*Proof.* Alg. 4 stops when  $r_1(P_i, \mathbf{P}_{i-1}^{k-1}) \leq r_1(P_{i-1}, \mathbf{P}_{i-1}^{k-1})$ . Using the definitions of  $r_1(P_i, \mathbf{P}_{i-1}^{k-1})$  and  $r_1(P_{i-1}, \mathbf{P}_{i-1}^{k-1})$ , 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 then be shown to be the solution of this inequality, using induction.  $\square$

*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.  $\square$

## 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 equilibrium finding in some cases) and various equilibrium properties under two different informational assumptions w.r.t. consumer preferences.

There are a number of directions remaining to be explored. The possibility of approximating best responses in the full information setting remains open. This problem doesn't 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 where the strategies are required to satisfy certain combinatorial constraints (e.g., matroid or knapsack) reflecting limits on individual catalogs would be of interest. Under such restrictions, our worst case PoA and PoS ratios might be improved. Connections to other game-theoretic models also bear exploration. For instance, allowing endogenous prices requires vendors to set prices, e.g., in a multi-vendor platform, offering a competitive extension of profit-maximizing, envy-free mechanisms (see e.g., (Guruswami et al. 2005)). Endogenous pricing has been considered in a recent competitive model related to ours, but where each vendor has a single item (Babaioff, Nisan, and Leme 2014).

## Acknowledgements

This research was supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada.

## References

Babaioff, M.; Nisan, N.; and Leme, R. P. 2014. Price competition in online combinatorial markets. In *Proceedings of the Twenty-Third International World Wide Web Conference (WWW-14)*, 711–722.

- Black, D. 1948. On the rationale of group decision-making. *Journal of Political Economy* 56(1):23–34.
- Chamberlin, J. R., and Courant, P. N. 1983. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review* 77(3):718–733.
- Feige, U. 1998. A threshold of  $\ln n$  for approximating set cover. *Journal of the ACM* 45(4):634–652.
- Gaertner, W. 2002. Domain restrictions. In *Handbook of Social Choice and Welfare*, Vol. 1, 131–170. Arrow, K. J.; Sen, A. K.; and Suzumura, K. (eds.). North Holland.
- Goldberg, A. V.; Hartline, J. D.; Karlin, A. R.; Saks, M.; and Wright, A. 2006. Competitive auctions. *Games and Economic Behavior* 55(2):242–269.
- Guruswami, V.; Hartline, J. D.; Karlin, A. R.; Kempe, D.; Kenyon, C.; and McSherry, F. 2005. On profit-maximizing envy-free pricing. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA-05)*, 1164–1173.
- Honhon, D.; Jonnalagedda, S.; and Pan, X. A. 2012. Optimal algorithms for assortment selection under ranking-based consumer choice models. *Manufacturing & Service Operations Management* 14(2):279–289.
- Kleinberg, J.; Papadimitriou, C.; and Raghavan, P. 2004. Segmentation problems. *Journal of the ACM* 51:263–280.
- Li, X.; Nukala, S.; and Mohebbi, S. 2013. Game theory methodology for optimizing retailers’ pricing and shelf-space allocation decisions on competing substitutable products. *International Journal of Advanced Manufacturing Technology* 68(1-4):375–389.
- Lu, T., and Boutilier, C. 2011. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI-11)*, 280–286.
- Mallows, C. L. 1957. Non-null ranking models. *Biometrika* 44:114–130.
- Marden, J. I. 1995. *Analyzing and Modeling Rank Data*. London: Chapman and Hall.
- Martínez-de Albéniz, V., and Roels, G. 2011. Competing for shelf space. *Production and Operations Management* 20(1):32–46.
- Monroe, B. L. 1995. Fully proportional representation. *American Political Science Review* 89(4):pp. 925–940.
- Schön, C. 2010. On the product line selection problem under attraction choice models of consumer behavior. *European Journal of Operational Research* 206(1):260–264.
- Skowron, P.; Faliszewski, P.; and Slinko, A. 2013. Fully proportional representation as resource allocation: Approximability results. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI-13)*, 353–359.