

# Stability, Fairness And The Pursuit of Happiness in Recommender Systems

Gerdus Benadè, Nachiketa Sahoo  
Boston University

Top- $k$  personalized recommendations are ubiquitous, but do they make stable matches? We study whether, given complete information, buyers and sellers would participate in matches formed by top- $k$  recommendations instead of pursuing alternative matches among themselves. When there are no constraints on the number of times an item is recommended, we observe that top- $k$  recommendations are stable. When exposures are constrained, e.g., due to limited inventory or exposure opportunities, stable recommendations need not exist. We show that maximizing total buyer welfare is closely related to stability. Maximizing welfare under unit exposure constraints is stable, Pareto optimal and swap-envy free for orthogonal buyers, identical buyers, and buyers with dichotomous valuations. Most of these properties remain under arbitrary exposure constraints. Finally, we evaluate variants of common strategies for recommending under exposure constraint and find that they leave room for substantial instability and envy in three real-world datasets. Among them, maximizing buyer welfare leads to the most stable outcomes and near-zero swap-envy.

---

## 1. Introduction

Recommender systems facilitate markets by matching buyers to products (and their sellers) in large online platforms. They do so by learning buyers' preferences from past ratings and recommending to each buyer a subset of products she would like, from which the buyer may choose one. Such recommendations focus buyers' attention and cut down the number of costly product evaluation they need to do. Traditionally, recommender systems have focused on satisfying buyers, with the implicit assumption that matching buyers to products they like also benefits sellers by increasing sales and attracting more buyers to the market.

Recent research, however, questions this assumption. Buyer-focused recommendations can concentrate sales on popular sellers and increase inequities (Fleder and Hosanagar 2009). This endangers marketplaces employing them, since disgruntled sellers may withdraw inventory and target buyers through off-platform channels. As a result, there are calls to design recommender systems that serve all stakeholders — buyers, sellers, and the platform (Abdollahpouri et al. 2020).

Such systems must exhibit multi-sided fairness while respecting different stakeholders' constraints. For example, recommending a physical good to more buyers than the number of copies a seller has can result in a costly stock-out. On platforms like ad-networks and retail-websites

that provide sponsored product placements (e.g., Amazon.com) exposure is a limited resource. Contracted exposure obligations limit how many times a seller can be recommended to potential buyers. Sellers, in turn, would like scarce exposures to target those buyers who give them the best chance of sales.

Against this backdrop, we take inspiration from the matching literature to study the stability of markets formed by recommender systems. In traditional matching markets (e.g., students-to-schools or interns-to-hospitals), each side has a preference over the limited resource available on the other (e.g., seats and students) which can lead to instability unless carefully matched (Roth and Sotomayor 1990). Recommendations, likewise, engage customers’ limited attention and consume sellers’ limited exposure opportunities. Two-sided preferences over these limited resources create incentives to find better recommendations.

Stability is important in traditional (offline) matching markets. A high cost of discovery and incomplete information can let the market operate for some time, but unstable markets eventually unravel due to persistent incentives (Roth 1984a, 2015). Discovery is easier with digital platforms. Competition among digital platforms facilitates multi-homing and other forms of off-platform transactions (Bryan and Gans 2019, Bakos and Halaburda 2020). As in offline markets, a digital platform relying on unstable matches risks its users migrating to alternatives.

In this context, the central question we study is

*“Are personalized recommendations stable?”*

More specifically, would buyers and sellers, given complete information on preferences and constraints on both sides, participate in a platform using a personalized recommender system? Or may some prefer to pursue off-platform matches and transactions among themselves? If these markets are unstable, how can we design personalized recommender systems to mitigate the risk of unravelling?

### 1.1. Our Contributions

We study a setting in which each of a batch of buyers is simultaneously recommended  $k$  items, and each item is subject to a constraint on the number of times it is recommended. The simultaneous presence of multiple buyers and sellers (thickness) and their interactions are essential properties of markets (Roth 2007). On large online platforms multiple buyers have active shopping sessions at any time, permitting batch recommendation. Meanwhile, cognitive capacities limit how many items buyers can consider at a time. Recommendation channels (web pages, emails, etc.), too, have limited usable space. Therefore, recommender systems present a small fixed number of (3–10) items to a buyer at a time. This is the type of recommendation we study.

We start with [McFadden \(1973\)](#)’s choice model for the behaviour of a buyer who is recommended a set of items (or choice set). Under this, a buyer’s utility for an item consists of two components. The first component can be estimated by the platform at the time of recommendation, for example, from the buyer’s prior interactions using a learning algorithm like a collaborative or content-based filter. The second, random, component is unknown to the platform and buyer at the time of recommendation and manifests after careful consideration of the product — a process which takes the buyer time and effort ([Shugan 1980](#), [Alba and Hutchinson 1987](#)).<sup>1</sup>

We investigate three properties towards robust recommendations: *stability*, *Pareto optimality* and a form of *envy-freeness*. Pareto optimality (PO) of buyer utilities ensures efficiency: increasing the utility of any buyer comes at the cost of another. Envy-freeness says that no buyer prefers the recommendations made to another over her own. Since envy-freeness cannot be guaranteed for indivisible goods, we propose a relaxation suitable for  $k$ -recommendation called swap-envy-freeness up to one good (SEF1) which allows for envy but only to the extent that it can be eliminated by exchanging a pair of recommended items. *Stability* requires that buyers and sellers would prefer to be matched per the recommendations of the system rather than make side-deals ([Roth and Sotomayor 1990](#)). We call a recommendation profile (the collection of recommendations to all buyers) stable if there is no buyer-seller pair who would both strictly benefit from being paired based on the *known* utility components. As a result, this (ex-ante) stability is a property of the platform’s recommendations.

We study the existence of stable recommendations in [Section 3](#). In the absence of exposure constraints, when each item can be recommended an arbitrary number of times, we observe that top- $k$  recommendations are stable, envy free, and Pareto optimal. Limiting the item exposures, however, may make it impossible to find a stable recommendation ([Theorem 1](#)). We show that any potential gain from deviating from the platform’s recommendations, for both buyers and sellers, is a result of disparities in buyers’ valuations of items and the values of their choice sets. In fact, the gain can be (tightly) upper bound by the product of two factors that measure these differences.

In [Section 4](#) we study the properties of buyer welfare maximizing recommendations and show it is intrinsically linked to stability. Maximizing welfare is provably stable when there are no  $k$ -constraints on the size of the recommendation set ([Theorem 3](#)). In [Sections 4.2](#) and [4.3](#) we study recommendations under unit exposure constraints in two restricted preference domains: identical preferences, as might be the case with commonly agreed item qualities, and dichotomous values,

<sup>1</sup> Often, the unknown “error” component stands for the component unobserved by the econometrician (see, for example, [Manski \(1977\)](#)). Our framing, in which even a buyer may not know her own precise utility beforehand due to lack of complete product knowledge and variations across consumption occasions is compatible with this view and is in line with what is known in consumer psychology literature (see, for example [Hauser and Wernerfelt \(1990\)](#)).

where each buyer deems each item to have either high or low value. In both settings, maximizing the total expected buyer utility leads to stable, Pareto optimal and SEF1  $k$ -recommendations. Notably, maximizing expected buyer utility under our choice model is equivalent here to maximizing the *Nash welfare* with respect to the exponentiated utilities, which has several attractive fairness properties (Caragiannis et al. 2019).

These results are extended to arbitrary upper bounds on the number of times an item is recommended, with some caveats. Stability is not guaranteed for identical buyers in general (Example 3), though maximizing the buyer welfare remains PO and SEF1 (Theorem 5). In the dichotomous value setting, swap-envy-freeness is not guaranteed (Example 4), but maximizing buyer welfare remains stable and PO (Theorem 7).

Absent the guarantee of stable  $k$ -recommendations for arbitrary preferences, we propose alternative benchmark recommendation algorithms under exposure constraints in Section 5. We show round robin recommendations are SEF1 under unit exposure constraints (Theorem 8). Greedy top- $k$  recommendations, which is suitable for online decision making, is not stable, PO or SEF1. Finally, we outline how any batch recommendation algorithm can be deployed in an online manner when users arrive and receive recommendations one at a time, rather than in batches.

We conclude with a computational study using datasets from Amazon and Rent-the-runway. The algorithms all result in unstable recommendations in which a large fraction of participants have some incentive to deviate from the platform’s recommendations. The magnitude of this incentive, as an improvement over the participant’s current utility, range from as high as 150% for greedy top- $k$  to only 8% for the welfare maximizing recommendations.

Our theoretical and computational results suggest that platforms worried about participants pursuing off-platform transactions should maximize buyer welfare. It is provably stable in some settings and, even when stable recommendations can not be guaranteed, leads to the lowest incentive to deviate from the platforms recommendations.

## 1.2. Related Work

This paper builds on ideas from recommender systems, matching, and fair division. We briefly discuss the closest work from each.

Multi-stakeholder recommendations, which consider multiple buyer, seller and platform preferences, have become increasingly popular (Burke et al. 2016, Nguyen et al. 2017, Abdollahpouri et al. 2020). Fair recommendation in this context considers multiple stakeholders’ objectives (Ekstrand et al. 2022, Abdollahpouri and Burke 2019, Patro et al. 2020). Fair allocation is difficult to achieve under one-shot recommendation, though proportional fairness may be achieved considering recommendations over time (Chakraborty et al. 2017). Bateni et al. (2022) propose a stochastic

approximation scheme based on the Eisenberg-Gale convex program for an online advertising system which maximizes platform revenue while being approximately fair towards buyers. Patro et al. (2020) propose a greedy version of round robin which is shown to be envy-free up to one good (EF1) for buyers and guarantees sellers some exposure. In the context of recommending bundles to groups that consume them together, Serbos et al. (2017) show that maximizing fairness towards *all* members of a group is NP-hard and offer greedy approximation algorithms.

Questions of fairness become salient when buyers or sellers face constraints. Recommendation under capacity or exposure constraint has been studied, generally to maximize sales (Makhijani et al. 2019) or user utility (Sürer et al. 2018)—rarely to unpack participation incentives. Studying incentives, Tennenholtz and Kurland (2019) point out that the standard relevance based ranking by content recommender systems can encourage homogeneous content. Related work offers Shapley value based probabilistic recommendations as a solution (Ben-Porat and Tennenholtz 2018). In one of the few papers in the recommendations literature that uses stable matches, Eskandarian and Mobasher (2020) propose a deferred acceptance algorithm to diversify recommendations. We add to this growing body of literature by going beyond fairness, by asking if the participants will have the incentive to participate in a match produced by recommender system, under a one-shot matching scenario. Drawing inspiration from the Eisenberg-Gale program, we find preference domains where the interests of buyers and sellers can be aligned.

There is a rich literature on stable matchings dating back to the 1950’s (Stalnaker 1953, Gale and Shapley 1962). Roth and Sotomayor (1990) and Abdulkadiroglu and Sönmez (2013) offer thorough treatments of the topic. Recommending  $k$  items to a buyer reminds of many-to-one and many-to-many matchings. In many-to-one worker-firm or college admissions matching workers (students) are matched to firms (colleges), sometimes subject to quotas on the number of matches. When preferences satisfy a substitutability condition and when workers’ preferences depend only on the firm they apply to, not their colleagues, stable matchings exist (Kelso and Crawford 1982, Echenique and Oviedo 2006).<sup>2</sup> In fact, in traditional firm-proposing deferred acceptance schemes used to find stable match (e.g., Roth (1984b)), a worker can safely reject all but the top proposing firm since her preference over firms is assumed not to depend on colleagues, which allows stability. This assumption does not hold for sellers when a buyer chooses one of  $k$  recommended items: the probability that an item is selected depends on the other  $k - 1$  items. Thus, in  $k$ -recommendations, a seller can not accurately judge the attractiveness of being recommended to a buyer until the other  $k - 1$  items in that buyer’s choice set are fixed.

The literature on matching with externalities is most relevant to us. Three stability notions of increasing strength, which Bando et al. (2016) call conservative, pairwise and optimistic stability,

<sup>2</sup> Informally, preferences are substitutable if each item chosen from a set is also chosen from a subset that includes it.

have been proposed. Sasaki and Toda (1996) show that conservative stability (the weakest of the three) can be guaranteed in one-to-one settings by using a standard deferred acceptance algorithm on transformed preferences. A similar argument guarantees conservative stability in our many-to-many setting. The bulk of the literature focuses on pairwise stability in various settings, including sufficient conditions for existence (Mumcu and Saglam 2010) and many-to-one matching with externalities over firms (Bando 2014, 2012, Echenique and Oviedo 2004). Many-to-one matching with preferences over colleagues is similar to recommendations in which sellers have preferences over their competition in a choice set. Dutta and Massó (1997) show that if preferences over firms and colleagues satisfy a fairly restrictive lexicographic ordering condition a pairwise stable match exists. We also study pairwise stability, and note this condition is unlikely to hold when seller-side preferences are induced by a choice model. Echenique and Yenmez (2007) offer algorithms to compute stable matches *when they exist* under preferences over colleagues. Pycia (2012) show that when there is substantial variability of preferences, stable coalitions can be guaranteed only when the common members in a pair of coalitions prefer the same coalition. More recently, Pycia and Yenmez (2022) study stability under externality over agents outside of the matched group and show existence under a specific monotone externality property.

In the fair division literature, envy-free allocations (Foley 1967) and relaxations thereof (Lipton et al. 2004) have been studied for divisible (Brams and Taylor 1995, Procaccia 2016) and indivisible goods (Alkan et al. 1991, Lipton et al. 2004, Caragiannis et al. 2019) in both static and dynamic settings (Benadè et al. 2018, Zeng and Psomas 2020, Benadè et al. 2022). The concept of Nash welfare, or the product of agent utilities, originated in John Nash’s solution to a bargaining problem (Nash 1950). Maximizing Nash welfare when allocating indivisible goods among agents with additive utilities is known to be EF1 and PO (Caragiannis et al. 2019). Maximizing Nash welfare is typically NP-hard; Caragiannis et al. (2019) propose a computational approach which scales to reasonably sized instances. The notions of balance and impartiality we use in Section 3 also appear in Huang et al. (2022).

Two-sided fairness has recently received attention in the fair division literature (Gollapudi et al. 2020, Freeman et al. 2021, Igarashi et al. 2022). Caragiannis and Narang (2022) independently propose envy-freeness up to a single exchange of items in a setting where goods and chores are repeatedly matched to agents and find that a variation of round robin allocation adapted for repeated matchings is both EF1 and SEF1. We make a similar observation for round robin allocations in Section 5.1 and further show that maximizing expected buyer welfare is also SEF1. In the context of matching teams to players, Igarashi et al. (2022) propose two stability notions. An allocation is *swap stable* when there are no two players in two different teams so that swapping the players makes at least one of the four parties better off while leaving none worse off. An allocation

is *individually stable* when no player can deviate to another team without making one of the teams involved worse off. In keeping with the original stable marriage problem (Gale and Shapley 1962) and the notion of pairwise stability as studied in the context of matching with externalities, we require only that the deviating buyer and seller are strictly better off, not that the other parties involved are no worse off.

## 2. Model Formulation

Let  $\mathcal{B}$  denote a set of  $n$  buyers and  $\mathcal{I}$  a set of  $m$  items. We assume that every item is sold by a different seller and occasionally blur the distinction between recommending an item and recommending a seller.

A ( $k$ -)recommendation to buyer  $b \in \mathcal{B}$  is a set of  $k$  unique items  $\bar{A}_b \subseteq \mathcal{I}$ ,  $|\bar{A}_b| = k$ . Additionally, buyer  $b$  has (fixed) outside option  $\omega_b$ , which represents not selecting any of the recommended items and instead sticking with the status quo, continuing searching, or selecting an item not available on the platform. Let  $A_b = \bar{A}_b \cup \{\omega_b\}$  be the choice set of buyer  $b$ . We call the vector of recommendations  $A = (A_b)_{b \in \mathcal{B}}$  a *recommendation profile*. The buyers who are recommended item  $i$  are denoted  $A_i^{-1} = \{b \in \mathcal{B} : i \in A_b\}$ .

Let  $\mathcal{A}$  denote the set of feasible recommendation profiles. A recommendation profile is feasible if it satisfies constraints on the number of exposures received by each item, encoded as  $|A_i^{-1}| \leq c_i$  for all  $i \in \mathcal{I}$ . We typically assume unit exposure constraints for illustrative purposes, i.e.  $c_i = 1$  for all  $i \in \mathcal{I}$ . Under unit exposures  $|A_i^{-1}| = 1$ , we overload notation to let  $A_i^{-1}$  also refer to the (unique) buyer recommended item  $i \in \mathcal{I}$ . We use  $i$  interchangeably with the singleton set  $\{i\}$ .

Buyer behavior is assumed to follow a discrete choice model (McFadden 1973).<sup>3</sup> Buyer  $b \in \mathcal{B}$  has utility  $U_{bi} = u_{bi} + \varepsilon_{bi}$  for item  $i \in \mathcal{I}$ , where  $u_{bi}$  is the expected utility that  $b$  has for  $i$  and  $\varepsilon_{bi}$ , drawn independently and identically from a Gumbel distribution, is an unknown random component. We take the expected value  $u_{bi}$  as arbitrary and known to the platform—it may have been estimated from prior interactions using a collaborative filter. The random component,  $\varepsilon_{bi}$ , is not known in advance and can only be realized by the buyer after careful consideration of the item, a process which may be time-consuming and cognitively demanding. We call  $v_{bi} = e^{u_{bi}}$  buyer  $b$ 's *virtual value* for item  $i$ . Together the buyers, items and expected utilities constitute an *instance* over which recommendations are made.

Outside options are modeled similarly: the known component of buyer  $b$ 's outside option  $\omega_b$  is  $u_{b\omega}$ . We assume all  $b \in \mathcal{B}$  have  $u_{b\omega} = u_\omega$ . This is largely without loss of generality: when buyers

<sup>3</sup> Choice models have been extensively used in the recommender systems literature to learn user-preferences from their selections (Song et al. 2019, Moins et al. 2020, Song et al. 2022) and to predict user-choices when presented with a set of recommendations (Fleder and Hosanagar 2007, Chen et al. 2019, Carroll et al. 2021).

have outside options with different values, we can create normalized utilities  $u'_{bi} = u_{bi} - u_{b\omega}$  and all our results remain true when phrased in terms of the normalized utilities. Most examples set  $u_{\omega} = 0$  ( $v_{\omega} = 1$ ) for concreteness, though any constant will do.

Per choice model theory, buyer  $b \in B$  considering choice set  $S$  uncovers the previously unknown random components of item utilities, potentially at some cost. Then the buyer (deterministically) selects the option  $i \in S$  that provides greatest utility. McFadden (1973) shows the resulting probability that option  $i$  will be selected from choice set  $S$  is  $\mathbb{P}(b, i, S) = e^{u_{bi}} / \sum_{j \in S} e^{u_{bj}} = v_{bi} / \sum_{j \in S} v_{bj}$ .<sup>4</sup> Note that when  $v_{\omega} \neq 0$  a buyer's probability of selecting the outside option is positive and decreases in the quality of their  $k$ -recommendation.

The welfare of (the seller of) item  $i$  is assumed only to be increasing in the expected number of times  $i$  is selected, denoted  $\mathbb{E}_i(A) = \sum_{b \in A_i^{-1}} \mathbb{P}(b, i, A_b)$ . Under unit exposures, this simplifies to  $\mathbb{E}_i(A) = \mathbb{P}(i, A) \triangleq \mathbb{P}(A_i^{-1}, i, A_{A_i^{-1}})$ . This is flexible enough to capture sellers having different profit margins and expected profit increasing with the expected number of sales.

At the time of recommendation, a buyer's *expected utility* from the choice set  $S$  is given by  $u_b(S) = \mathbb{E}(\max_{i \in S}(U_{bi})) = \log(\sum_{i \in S} e^{u_{bi}})$  (Williams 1977).<sup>5</sup> The *utility profile* associated with recommendation profile  $A$  is  $u(A) = (u_b(A_b))_{b \in B}$ . Accordingly,  $v_b(S) = e^{u_b(S)} = \sum_{i \in S} v_{bi}$  and we call  $v(A) = (v_b(A_b))_{b \in B}$  the *virtual value profile*. We refer to  $\sum_{b \in B} u_b(A_b)$  as the (total) *buyer welfare*.

Notice that each buyer has a fixed preference order over  $\mathcal{I} \cup \{\omega\}$ , as determined by  $\{u_{bi} : \forall i \in \mathcal{I} \cup \{\omega\}\}$ . However, sellers of items do not have a fixed preference order over buyers: the probability of being selected depends both on the item's utility and on how much competition the item faces in a given choice set (specifically, the sum of virtual values of other options in the choice set).

## 2.1. Measuring The Quality of a Recommendation

*Efficiency* A natural requirement is that the recommendation profile is *Pareto optimal* (PO) with respect to buyer utilities. Let  $[n] = \{1, \dots, n\}$ . A vector  $x \in \mathbb{R}^n$  *strictly dominates*  $y \in \mathbb{R}^n$  when  $x_i \geq y_i$  for all  $i \in [n]$  and there exists  $j \in [n]$  where  $x_j > y_j$ . A recommendation profile  $A \in \mathcal{A}$  is Pareto optimal if there does not exist another recommendation profile  $A' \in \mathcal{A}$  such that  $u(A')$  strictly dominates  $u(A)$ . Notice that  $u(A)$  is undominated exactly when  $v(A)$  is undominated.

*Fairness* A standard notion of fairness is *envy-freeness*, which we consider from the buyer perspective. Envy-freeness requires that each buyer prefers her choice set over that of any other. Formally, recommendation profile  $A$  is envy-free when  $u_b(A_b) \geq u_b(A_{b'})$  (equivalently,  $v_b(A_b) \geq$

<sup>4</sup> This expression for probability as a function of  $u_{bi}$ , incidentally, is the softmax function that is widely used during training of neural networks with discrete outcome variables (Goodfellow et al. 2016, p. 81).

<sup>5</sup> Expected utility of choice sets thus calculated has been used to measure consumer welfare for policy evaluation. Train (2009, Ch. 3.5) provides a textbook discussion; De Jong et al. (2007) surveys theory and real-world applications.



$v_b(A_{b'})$ ) for all  $b, b' \in \mathcal{B}$ . Envy-freeness is often impossible with indivisible objects (consider allocating a single valuable item to two agents). Thus it is commonly relaxed to envy-freeness up to one item (EF1), which allows envy to exist but only to the extent that it can be eliminated by removing a single item from the envied agent's allocation. In our setting, each buyer must be recommended exactly  $k$  items and simply removing an item from a buyer's choice set is not an option. We propose a relaxation of envy-freeness, called *swap-envy-freeness*, to accommodate this. A recommendation profile is *swap-envy-free up to 1 item* (SEF1) when any pairwise envy between buyers  $b \neq b' \in \mathcal{B}$  can be eliminated by exchanging a single pair of items between them. Formally,  $A$  is SEF1 if, for all  $b, b' \in \mathcal{B}$  where  $b$  envies  $b'$ , there exist a pair of items  $i \in A_b, j \in A_{b'}$  so that  $u_b(A_b \cup j \setminus i) \geq u_b(A_{b'} \cup i \setminus j)$ .

*Ex-ante stability* We combine the pairwise stability notion in Bando et al. (2016) with the preference orders induced by the expected utilities under the choice model. A buyer-item pair  $(b, i) \in \mathcal{B} \times \mathcal{I}$  is called a *blocking pair* in recommendation profile  $A$  if both the buyer and the item's seller strictly benefit in the recommendation profile which results from  $i$  joining  $b$  and restoring feasibility. Formally, for unit exposure constraints,  $(b, i)$  blocks when  $i \notin A_b$  and there exists a item  $j \in A_b$  so that both  $u_{bi} > u_{bj}$  and  $\mathbb{P}(i, A) < \mathbb{P}(b, i, A_b \cup i \setminus j)$ . A recommendation profile is *stable* in the absence of a blocking pair.

Restoring the feasibility of a  $k$ -recommendation after  $i \in A_c$  replaces some item  $j$  in  $A_b$  requires that there is an item available that can feasibly be recommended  $c$ , the buyer  $i$  left, since each buyer must be recommended  $k$  items. Under unit exposures, recommending  $j$  to  $c$  is feasible since we are guaranteed  $j \notin A_c$ . To handle general exposure constraints, where possibly  $j \in A_c$  (meaning that the ejected item can not be recommended to  $c$ ) we assume there are dummy items available for which every buyer has the lowest possible value which may be used to restore feasibility. When there are no restrictions on the number of items recommended to each buyer simply including  $i$  in  $A_b$  is feasible, there is no need to restore  $c$ 's choice set to size  $k$ .

A benefit of studying ex ante stability entirely in terms of buyers' expected utilities  $u_{bi}$  and sellers' expected sale  $\mathbb{P}(i, A)$ , as we do here, is that it is a property of a recommendation profile within the control of the platform. It is also in line with the literature where stability is typically defined using ex-ante preferences available to the central planner<sup>6</sup> — this translates to using only the known component in a random utility model.

One may be tempted, instead, to define (an ex-post) stability based on the item utilities ( $U_{bi} = u_{bi} + \varepsilon_{bi}$ ). Immediately, a platform will be unable to tell if a fixed recommendation profile is stable or not, since it depends on the realizations of the random  $\varepsilon_{bi}$  components. This is also impractical since

<sup>6</sup> For example, doctors' (hospitals') preferences over hospitals (doctors) before they are matched (see Roth 1984a, pp. 996, 1008)

it requires buyers to know their precise utilities for potentially thousands of items. Moreover, the problem collapses to effectively assigning one item to each buyer — each buyer will deterministically choose the item with highest (realized) utility and such (ex-post) stable matches trivially exist.

Finally, it is worth distinguishing between instability and choosing the outside option  $\omega$  from a choice set. When a buyer selects her outside option  $\omega$  (prefers it to items she was recommended) the platform loses a transaction. It is still possible that she is satisfied with the *way she was matched*, i.e., no better match that sellers would agree to was available (the match was stable). Conversely, she may, for the time being, select one of the recommended items in an unstable match. However, over time the incentives present may lead users to multi-home or migrate to rival platforms. To drive home this distinction, all examples remain valid when ignoring the outside option (set  $v_\omega = 0$ ).

We illustrate Pareto optimality, envy-freeness and stability with an example.

EXAMPLE 1. Consider an instance with buyers  $\{1,2\}$ , four items  $\{a,b,c,d\}$ , outside option  $\omega$  with  $u_\omega = 0$  and  $k = 2$ . Table 1 shows the virtual values for an instance with identical buyers and two good and two bad items. Suppose each item can only be recommended once.

**Table 1** Buyers' virtual values.

	$a$	$b$	$c$	$d$	$\omega$
Buyer 1	2	2	1	1	1
Buyer 2	2	2	1	1	1

The outside option appears in every buyer's choice set, we omit explicitly including it except where an omission could create confusion. The recommendation profile  $(\{a,b\}, \{c,d\})$  is Pareto optimal even though it recommends both good items to the same buyer. It is also SEF1: buyer 2 envies buyer 1 but swapping  $b$  with  $c$  removes the envy. Furthermore, it is unstable with blocking pair  $(2,b)$ : 2 prefers  $b$  over both  $c$  and  $d$ , and  $b$  has a larger probability of being purchased by 2 in the choice set  $\{b,c\}$  than by 1 in the choice set  $\{a,b\}$ . Recommendation profile  $(\{a,c\}, \{b,d\})$  is Pareto optimal, envy free and stable.

### 3. On The Existence of Stable Recommendations

We investigate the conditions under which stable recommendations exist. When the exposure constraints do not prevent recommending each buyer her  $k$  most preferred items, top- $k$  recommendations are stable. Omitted proofs appear in the appendix.

PROPOSITION 1. *In any instance where the unconstrained top- $k$  recommendation profile is feasible, for example, when  $\vec{c} = \infty$  or when  $\vec{c} = 1$  and all buyers have disjoint sets of top- $k$  items, recommending each buyer their  $k$  most liked items is feasible, stable, welfare optimal and envy-free.*

Stability, here, follows from the fact that a buyer recommended her  $k$  highest value items will not participate in a blocking pair. Recommender systems are often deployed on digital products where exposure constraints are less likely to exist. For example, there may be no limit on the number of times a streaming platform can stream a particular movie, or a book seller may sell an e-book. It is reassuring to know that personalized top- $k$  recommendation is stable in these settings. To the best of our knowledge, this property of top- $k$  recommendation has not been discussed before.

Unfortunately, once the exposure constraints make it impossible to recommend to every buyer their most preferred  $k$  items, stable recommendations need no longer exist.

**THEOREM 1.** *Under unit exposures, there exist instances where no stable recommendation exists.*

*Proof.* Consider the instance in Table 2 with buyers  $\{1, 2\}$  and sellers  $\{a, b, c, d\}$ . We argue that this instance with  $k = 2$  does not permit a stable recommendation under unit exposure constraints.

**Table 2** Buyers' virtual values.

	$a$	$b$	$c$	$d$	$\omega$
Buyer 1	10	1	7	6	1
Buyer 2	10	8	4	5	1

Since  $A_2 = S \setminus A_1$ , we need only check all possible  $A_1$ . For  $A_1 \in \{\{ac\}, \{ad\}\}$ ,  $(2, a)$  is a blocking pair as buyer 2 prefers  $a$  over all their recommended sellers and would be willing to eject seller  $b$  from their current choice set. For  $A_1 \in \{\{ab\}, \{bc\}, \{bd\}\}$ ,  $(2, b)$  blocks, since 2 will always be willing to accept  $b$  (it is one of 2's top-2 items), and  $b$ 's purchase probability, which is at most  $1/8$  when it is recommended to buyer 1, increases. Finally, for  $A_1 = \{c, d\}$ ,  $(1, a)$  blocks since  $a$  is 1's most liked item and the purchase probability of  $a$  is higher when competing with  $c$  or  $d$  in 1's choice set than when competing with  $b$  in 2's choice set.  $\square$

It is instructive to consider a few variations of the example in Table 2. First, setting  $v_\omega = 0$  does not lead to stability. In other words, the presence of outside options is not the source of instability. Instead, it comes from the externalities sellers have over other sellers they are recommended alongside. Consider a variation where a buyer buys each item in their choice set independently with probability increasing in the expected utility of the item (i.e., a buyer can buy multiple items). Now seller preferences are independent of the choice set, buyer-preferences remain substitutable ( $u_b(S) = \sum_{i \in S} u_{bi}$ ), and existing work (say, Kelso and Crawford (1982)) imply stability even in the presence of buyer and item side constraints —  $\{\{ac\}, \{bd\}\}$  and  $\{\{cd\}, \{ab\}\}$  are stable.

A closer look at the instance in Theorem 1 reveals two situations which facilitate blocking pairs. First, whenever  $b$  is recommended to buyer 1, the massive discrepancy in the buyers' values for

item  $b$  created the potential for  $b$  to increase its purchase probability by deviating to 2. Second, the allocations with  $b \in A_2$  lead to a large disparity in the bundle values, specifically,  $|v_1(A_1) - v_2(A_2)| \geq 4$ . When choice sets have very different values it leads to unequal levels of competition, and even an item valued identically by the buyers may benefit from deviating to seller where it will face less competition, as shown by the blocking pair  $(2, a)$  in allocation  $A_1 = \{a, c\}$ ,  $A_2 = \{b, d\}$ . We can formalize this intuition that the incentive to participate in a blocking pair depends on the difference in the buyers' valuations of the same item and the difference in competition across choice sets in terms of *balance* and *impartiality*, respectively (Huang et al. 2022).

**DEFINITION 1 ( $\alpha$ -BALANCE).** The valuations of an item  $i$  is called  $\alpha_i$ -balanced if  $v_{bi} \leq \alpha_i \cdot v_{ci}$  for all  $b, c \in \mathcal{B}$ . When each  $i \in \mathcal{I}$  is  $\alpha_i$ -balanced, the instance is called  $\alpha$ -balanced, with  $\alpha = \max_i \alpha_i$ .

When buyers have identical values the instance is 1-balanced.

**DEFINITION 2 ( $\beta$ -IMPARTIALITY).** An allocation  $A$  is called  $\beta$ -impartial when  $\beta$  is the smallest value such that  $v_b(A_b) \leq \beta \cdot v_c(A_c)$  for all  $b, c \in \mathcal{B}$ .<sup>7</sup>

For an arbitrary allocation  $A$  with blocking pair  $(b, i) \notin A$ , let  $A'$  denote an allocation with  $A'_b = A_b \cup i \setminus j$  for some  $j \in A_b$ . For example, under unit constraints if  $i \in A_c$  then  $A'$  can be identical to  $A$  except that  $A'_b = A_b \cup i \setminus j$  and  $A'_c = A_c \cup j \setminus i$ , i.e., the allocation that results from  $A$  when  $b$  deviates with  $i$  and the item ejected from  $A_b$  is recommended to  $c$ . We can upper bound the benefit from participating in a blocking pair in terms of the balancedness of the instance and the impartiality of  $A$ .

**THEOREM 2.** *Consider an  $\alpha$ -balanced instance and  $\beta$ -impartial allocation  $A$  with  $i \in A_c$ . For any blocking pair  $(b, i)$  of  $A$ , the multiplicative gain of buyer  $b$  and seller  $i$  when deviating from  $A$  to  $A'$  (as defined above) can be tightly upper bound as*

$$\frac{\mathbb{P}(i, A')}{\mathbb{P}(i, A)} \leq \alpha_i \beta \leq \alpha \beta, \quad \text{and} \quad \frac{v_b(A'_b)}{v_b(A_b)} < \alpha_i \beta \leq \alpha \beta.$$

**Proof.** We first bound the gain seller  $i$  can get from deviating. Observe that  $v_{bj} \leq v_{bi}$  since  $b$  ejected  $j$  in favor of  $i$  when deviating. Now

$$\begin{aligned} \frac{\mathbb{P}(i, A')}{\mathbb{P}(i, A)} &= \frac{v_{bi}}{v_b(A_b) + v_{bi} - v_{bj}} \bigg/ \frac{v_{ci}}{v_c(A_c)} \\ &= \frac{v_{bi}}{v_{ci}} \cdot \frac{v_c(A_c)}{v_b(A_b) + v_{bi} - v_{bj}} \\ &\leq \frac{v_{bi}}{v_{ci}} \cdot \frac{v_c(A_c)}{v_b(A_b)} \leq \alpha_i \beta. \end{aligned}$$

<sup>7</sup>  $\beta_b$ -impartiality can be defined with  $\beta_b = \max_{c \in \mathcal{B}} \{v_c(A_c)/v_b(A_b)\}$  and will lead to slightly stronger bounds.

We now bound the welfare increase of buyer  $b$ . Since  $(b, i)$  is a blocking pair the purchase probability of  $i$  increases after deviation, so

$$\mathbb{P}(i, A) < \mathbb{P}(i, A') \iff \frac{v_{ci}}{v_c(A_c)} < \frac{v_{bi}}{v(A_b) + v_{bi} - v_{bj}} \iff v_b(A_b) + v_{bi} - v_{bj} < \frac{v_{bi} \cdot v_c(A_c)}{v_{ci}}. \quad (1)$$

We can now bound buyer  $b$ 's multiplicative increase in welfare as

$$\frac{v_b(A'_b)}{v_b(A_b)} = \frac{v_b(A_b) + v_{bi} - v_{bj}}{v_b(A_b)} < \frac{v_{bi}}{v_{ci}} \cdot \frac{v_c(A_c)}{v_b(A_b)} \leq \alpha_i \beta,$$

where the second transition uses Equation (1).

We show in the Section EC.1 there exists instances where both the above bounds are tight.  $\square$

As an example, consider allocation  $A_1 = \{a, c, \omega\}$ ,  $A_2 = \{b, d, \omega\}$  in the instance of Theorem 1 which has  $\alpha_a = 1$  and  $\beta = \frac{18}{14}$ , implying a maximum benefit of  $\frac{18}{14} \approx 1.29$ . Participating in the blocking pair  $(2, a)$  increases seller a's purchase probability by a factor of  $\frac{10/16}{10/18} \approx 1.11$  when ejecting  $b$ . Buyer 2's welfare increases by a factor of  $\frac{16}{14} \approx 1.15$ .

When the unconstrained top- $k$  recommendation profile is feasible (as in Proposition 1), top- $k$  recommendations are stable and maximize buyer welfare. Next, we show that maximizing buyer welfare is closely related to stability.

## 4. Maximizing Buyer Welfare

A common platform objective is to maximize the total buyer welfare. Let  $A^*$  be a recommendation profile which maximizes total buyer welfare, i.e., under unit exposures  $A^*$  is an optimal solution to

$$\begin{aligned} & \text{maximize} && \sum_{b \in \mathcal{B}} \log \left( \sum_{i \in \mathcal{I}} e^{\hat{u}_{bi}} x_{bi} + e^{u_\omega} \right) \\ & \text{s.t.} && \sum_{b \in \mathcal{B}} x_{bi} = 1, \text{ for all } i \in \mathcal{I} && \text{(unit exposures)} \\ & && \sum_i x_{bi} = k, \text{ for all } b \in \mathcal{B}, \text{ and} && \text{(} k\text{-recommendations)} \\ & && x_{bi} \in \{0, 1\}, \text{ for all } b \in \mathcal{B}, i \in \mathcal{I}, \end{aligned}$$

where  $x_{bi} = 1$  when  $i$  appears in  $b$ 's choice set. This is equivalent to maximizing the product of buyer's virtual values, or maximizing the *Nash welfare* with respect to virtual values  $v_{bi} = e^{\hat{u}_{bi}}$ . In the fair division literature, maximizing Nash welfare (on item values) is known to be EF1 and PO (Caragiannis et al. 2019).

Because buyer welfare is directly maximized, we may expect buyers to be less willing to deviate from the platform's recommendations. However, even for unit exposures, maximizing buyer welfare is not guaranteed to find a stable recommendation in instances where one exists.

**Table 3** Maximizing buyers' welfare need not result in a stable recommendation even when one exists.

	$a$	$b$	$c$	$d$	$\omega$
Buyer 1	$\boxed{10}$	$\underline{6}$	$\boxed{3}$	1	$\underline{1}$
Buyer 2	10	$\boxed{9.5}$	0.5	$\boxed{0.25}$	$\underline{1}$

EXAMPLE 2. Consider the instance with virtual values as shown in Table 3. The recommendation that maximizes buyer welfare is  $A_1 = \{a, c\}, A_2 = \{b, d\}$  (boxed), however,  $(2, a)$  is a blocking pair. At the same time, the underlined recommendation  $A_1 = \{b, c\}, A_2 = \{a, d\}$  is stable.

Despite these limitations, we show in the remainder of the section that maximizing buyer welfare is closely tied to stability and leads to fair, efficient and stable recommendations in restricted preference domains.

#### 4.1. Recommendations without $k$ -constraints

Proposition 1 suggests that exposure constraints are obstacles to stability: in the absence of exposure constraints top- $k$  recommendations are stable. We now investigate the  $k$ -recommendation constraints and show that in their absence, when a buyer can be recommended an arbitrary number of items, maximizing buyer welfare is stable and efficient.

THEOREM 3. *Maximizing buyer welfare without  $k$ -constraints is stable and Pareto optimal.*

Let  $A^*$  be the welfare maximizing allocation. We first show  $A^*$  is stable. Suppose for contradiction  $i \in A_c^*$  and there is a blocking pair  $(b, i) \in \mathcal{B} \times \mathcal{I}$ . Since buyer  $b$  does not benefit from being recommended multiple copies of the same item, we conclude that  $i \notin A_b^*$ . Let  $A'$  be the recommendation profile which coincides with  $A^*$ , except that  $A'_b = A_b^* \cup \{i\}$  and  $A'_c = A_c^* \setminus \{i\}$ . In words,  $A'$  is the allocation that results from  $A^*$  after including  $i$  in  $b$ 's choice set. Note that, since there are no constraints on the number of items recommended to each buyer, there is no need to displace a seller from  $A_b^*$  to restore feasibility,  $A'$  is feasible.

Since  $(b, i)$  blocks, both buyer and seller strictly benefit from  $i$  being recommended to  $b$  rather than  $c$ . Buyers have non-zero virtual values for all sellers, so the buyer-side improvement is guaranteed. On the seller side, we conclude that

$$\mathbb{P}(b, i, A'_b) = \frac{v_{bi}}{v_b(A_b^*) + v_{bi}} > \frac{v_{ci}}{v_c(A_c^*)} = \mathbb{P}(c, i, A_c^*) \Leftrightarrow 1 - \mathbb{P}(c, i, A_c^*) > 1 - \mathbb{P}(b, i, A'_b) \quad (2)$$

where  $i \in A_c^*$ . Now

$$\begin{aligned} \frac{v_b(A'_b) \cdot v_c(A'_c)}{v_b(A_b^*) \cdot v_c(A_c^*)} &= \frac{(v_b(A_b^*) + v_{bi}) \cdot (v_c(A_c^*) - v_{ci})}{v_b(A_b^*) \cdot v_c(A_c^*)} \\ &= \frac{v_{bi}}{\mathbb{P}(b, i, A'_b)} \cdot \left( \frac{v_{ci}}{\mathbb{P}(c, i, A_c^*)} - v_{ci} \right) \Big/ \left[ \left( \frac{v_{bi}}{\mathbb{P}(b, i, A'_b)} - v_{bi} \right) \cdot \frac{v_{ci}}{\mathbb{P}(c, i, A_c^*)} \right] \end{aligned}$$

$$\begin{aligned}
&= \frac{v_{bi}}{\mathbb{P}(b, i, A'_b)} \cdot \frac{v_{ci} - v_{ci}\mathbb{P}(c, i, A_c^*)}{\mathbb{P}(c, i, A_c^*)} \bigg/ \left[ \frac{v_{bi} - v_{bi}\mathbb{P}(b, i, A'_b)}{\mathbb{P}(b, i, A'_b)} \cdot \frac{v_{ci}}{\mathbb{P}(c, i, A_c^*)} \right] \\
&= \frac{1 - \mathbb{P}(c, i, A_c^*)}{1 - \mathbb{P}(b, i, A'_b)} > 1,
\end{aligned}$$

where the final transition follows from Equation (2). This contradicts  $A^*$  being the welfare maximizing allocation; we conclude that the welfare maximizing allocation is stable.

Pareto optimality follows directly from  $A^*$  maximizing the product of utilities.  $\square$

This establishes that, although stability cannot be guaranteed in the presence of  $k$ -constraints, maximizing buyer welfare is intrinsically linked to stability. Interestingly, the same holds when integrality constraints are relaxed: then the buyer welfare maximization without the  $k$ -constraint is the Eisenberg–Gale convex program (Eisenberg and Gale 1959) and stability of its optimal solution (if integral) follows from the KKT conditions. The market clearing price,  $p_i$ , of item  $i$  is exactly the same as the purchase probability  $\frac{v_{bi}}{v_b(A_b^*)}$  from those buyers  $b$  who are recommended  $i$ . Buyers not recommended an item are willing to pay less, i.e., offer lower purchase probabilities. Similarly, each buyer receives her most valued bundle. Therefore no matched buyer and seller would block. A similar property holds when the optimal solution is fractional.

We next identify two restricted preference domains—buyers with identical or dichotomous values—where maximizing buyer welfare is attractive even in the presence of  $k$ -constraints.

## 4.2. Identical Buyers

Suppose the buyers agree on a common evaluation  $u_i$  of each item  $i \in \mathcal{I}$ . So,  $u_i = u_{bi}$  for all  $b \in \mathcal{B}$ . We first consider maximizing expected buyer welfare under unit exposure constraints.

**THEOREM 4.** *For unit exposure constraints and identical preferences,  $A^*$  is stable, PO and SEF1.*

*Proof.* We first show  $A^*$  is stable. Assume for contradiction it is not, then there exists a blocking pair  $(b, i) \in \mathcal{B} \times \mathcal{I}$ . Let  $c$  be the buyer currently recommended  $i$ . By definition, a blocking pair implies  $i \notin A_b^*$  and  $\exists j \in A_b^*$  so that  $v_{bj} < v_{bi}$  and  $\mathbb{P}(i, A^*) < \mathbb{P}(i, A')$ , where  $A'$  is constructed from  $A^*$  by exchanging  $i$  and  $j$ , i.e.  $A'_b = A_b^* \cup i \setminus j$  and  $A'_c = A_c^* \setminus i \cup j$  and  $A'_d = A_d^* \forall d \notin \{b, c\}$ .

Since  $A^*$  maximizes the product of virtual values  $v_b(A_b^*) \cdot v_c(A_c^*) \geq v_b(A'_b) \cdot v_c(A'_c)$ . It follows that  $\frac{v_i - v_j}{v_b(A'_b)} = \frac{v_b(A_b^*) + v_i - v_j - v_b(A_b^*)}{v_b(A'_b)} = 1 - \frac{v_b(A_b^*)}{v_b(A'_b)} \leq 1 - \frac{v_c(A'_c)}{v_c(A_c^*)} = \frac{v_c(A_c^*) - (v_c(A_c^*) - v_i + v_j)}{v_c(A_c^*)} = \frac{v_i - v_j}{v_c(A_c^*)}$ , from which we get  $v_b(A'_b) \geq v_c(A_c^*)$ . As a result,  $\mathbb{P}(i, A^*) = v_i/v_c(A_c^*) \geq v_i/v_b(A'_b) = \mathbb{P}(i, A')$ , contradicting the fact that  $i$  strictly benefited from participating in the blocking pair. We conclude no blocking pair exists and  $A^*$  is stable.

Finally, we show that  $A^*$  is SEF1 with respect to the virtual values, which implies SEF1 for the true utilities by the monotonicity of the logarithm. Assume for contradiction that  $A^*$  is not SEF1

for virtual values and suppose that buyer  $b \in \mathcal{B}$  envies  $c \in \mathcal{B}$ . By definition,  $v(A_b^*) < v(A_c^*)$  and, for every  $i \in A_c^*, j \in A_b^*$ ,  $v(A_b^* \setminus j \cup i) < v(A_c^* \setminus i \cup j)$ .

There exists at least one pair of sellers  $(j, i) \in (A_b^*, A_c^*)$  such that  $v_j < v_i$ . Otherwise,  $v_j \geq v_i$  for all  $j \in A_b^*, i \in A_c^*$ . Since  $|A_b^*| = |A_c^*|$ , it follows that  $v(A_b^*) \geq v(A_c^*)$ , which is not the case. Let  $A'$  be the recommendation that results from swapping  $i$  and  $j$ , in other words,  $A'_b = A_b^* \cup i \setminus j$ ,  $A'_c = A_c^* \setminus i \cup j$  and  $A'_d = A_d^*$  for  $d \in \mathcal{B} \setminus \{b, c\}$ .

Set  $\delta = v_i - v_j > 0$ , so  $v(A'_b) = v(A_b^*) + \delta$  and  $v(A'_c) = v(A_c^*) - \delta$ . Now

$$v(A'_b) \cdot v(A'_c) = (v(A_b^*) + \delta)(v(A_c^*) - \delta) = v(A_b^*) \cdot v(A_c^*) + \delta(v(A_c^*) - (v(A_b^*) + \delta)) > v(A_b^*) \cdot v(A_c^*),$$

where the final transition follows from

$$v(A_c^*) - (v(A_b^*) + \delta) = (v(A'_c) + \delta) - (v(A_b^*) + \delta) = v(A'_c) - v(A'_b) + \delta > 0$$

since  $A^*$  is not SEF1 by assumption and  $\delta > 0$ . This contradicts the fact that  $A^*$  maximized the product of virtual values. We conclude that  $A^*$  is SEF1.

Pareto optimality follows directly from  $A^*$  maximizing the product of buyers' virtual values.  $\square$

If it were possible to guarantee exactly equal values across all choice sets (i.e. find a perfectly impartial recommendation profile with  $\beta = 1$ ), then Theorem 2 would imply that the allocation is stable, since  $\alpha = 1$  for agents with identical preferences. Perfectly balancing choice set virtual values is too much to hope for, in fact, maximizing the product of buyer virtual values can lead to arbitrarily large values of  $\beta$  even when impartial allocations exist under general valuations. The previous result shows that, at least with identical preferences, it leads to choice sets with very similar utilities and gets close enough to ensure stability. When items can be recommended more than once, carefully constructed upper bounds may enforce less impartial recommendation profiles where stability can no longer be guaranteed. The following example illustrates this.

**EXAMPLE 3.** We study an instance with  $n = 2, k = 3$  and six unique items and values and upper bounds on the number of exposures as in Table 4. Without loss of generality, assume buyer 1 receives the most valuable item (the other case is symmetric). The resulting recommendation which maximizes the product of buyers' virtual values is indicated with checkmarks.

This welfare maximizing recommendation profile is not stable: consider transferring  $a$  to 2, displacing  $b$  from 2's current choice set. The only resulting recommendation is boxed (1 can not be recommended a second copy of  $b$ ). Buyer 2's total virtual value increased from 20 to 22. Item  $a$ 's probability of being selected also increased, from  $\frac{12}{26}$  to  $\frac{12}{22}$ . We conclude that  $(2, a)$  is a blocking pair.

Notice that in the corresponding instance with seven items (two distinct items with value 10) and unit constraints the product maximizing choice sets are valued more similarly, with virtual values 22 and 24 ( $\beta = \frac{24}{22} \approx 1.1$ , rather than 20 and 26 ( $\beta = \frac{26}{20} = 1.3$ )).



**Table 4** Buyers' virtual values.

	$a$	$b$	$c$	$d$	$e$	$f$	$\omega$
Values	12	10	5	4	3	3	1
Exposure limit	1	2	1	1	1	1	–
Buyer 1	✓	☑			☑	☐	☑
Buyer 2	☐	✓	☑	☑			☑

Despite stability breaking down,  $A^*$  remains PO (trivially) and SEF1. The argument that SEF1 holds is similar to the one presented in the proof of Theorem 4.

**THEOREM 5.**  $A^*$  is PO and SEF1 for buyers with identical preferences.

### 4.3. Dichotomous Values

We now consider the case where buyers have dichotomous preferences with either high or low value for every item (and are indifferent among items of each type). Formally, assume  $v_{bi} \in \{\ell, h\}$  for all  $b \in \mathcal{B}$  and  $i \in \mathcal{I}$ , for some real-valued  $\ell < h$ . Buyers need not agree on the evaluation of an item: it may be of high value to some and of low value to others. Here, maximizing the total buyer welfare (or maximizing Nash welfare with respect to the virtual values) again leads to stability.

**THEOREM 6.** For buyers with dichotomous values,  $A^*$  is stable, PO and SEF1 under unit exposure constraints.

*Proof of Theorem 6. Stability:* Assume, for contradiction,  $A^*$  is not stable. Then there exists a blocking pair  $(b, i) \in \mathcal{B} \times \mathcal{I}$ . Let  $c$  be the buyer to whom  $i$  is currently recommended.

Observe that  $b$ 's choice set contains at least one low-valued seller, otherwise she can not benefit from participating in a blocking pair. Let  $j \in A_b^*$  be such an item with  $v_{bj} = \ell$ . Construct  $A'$  from  $A^*$  by exchanging  $i$  and  $j$ . So  $A'_b = A_b^* \cup i \setminus j$ ,  $A'_c = A_c^* \cup j \setminus i$  and  $A'_d = A_d^*$  for all  $d \in \mathcal{B} \setminus \{b, c\}$ .

We now study the possible values of  $(v_{bi}, v_{ci})$ :

1.  $(\ell, \ell)$ : Participating in a blocking pair requires displacing a buyer with strictly lower value, which is not possible since  $v_{bi} = \ell$ . This contradicts  $(b, i)$  being a blocking pair.
2.  $(\ell, h)$ : As in the previous case,  $v_{bi} = \ell$ , which contradicts  $(b, i)$  being a blocking pair.
3.  $(h, \ell)$ : In this case  $i$  is assigned to  $c$  despite  $c$  having low and  $b$  having high value for it. It follows that  $v_b(A'_b) = v_b(A_b^*) - \ell + h > v_b(A_b^*)$  and  $v_c(A'_c) = v_c(A_c^*) - \ell + v_c(j) \geq v_c(A_c^*)$ . However, this implies  $v_b(A'_b) \cdot v_c(A'_c) > v_b(A_b^*) \cdot v_c(A_c^*)$ , contradicting the fact that  $A^*$  maximizes the product of virtual values.

4.  $(h, h)$ : Suppose that  $v_{cj} = h$ . Now  $v_b(A'_b) > v_b(A_b^*)$  and  $v_c(A'_c) = v_c(A_c^*)$ , contradicting that  $A^*$  maximizes the product of virtual values. We conclude that  $v_{cj} = \ell$ , so  $b$  and  $c$  value the swapped items identically. We can now argue as in Theorem 4 that

$$\begin{aligned} \frac{v_b(A_b^*) \cdot v_c(A_c^*)}{v_b(A'_b) \cdot v_c(A'_c)} &\geq 1 \\ 1 - \frac{v_b(A_b^*)}{v_b(A'_b)} &\leq 1 - \frac{v_c(A'_c)}{v_c(A_c^*)} \\ \frac{(v_b(A_b^*) + v_{bi} - v_{bj}) - v_b(A_b^*)}{v_b(A'_b)} &\leq \frac{v_c(A_c^*) - (v_c(A_c^*) - v_{ci} + v_{cj})}{v_c(A_c^*)} \\ \frac{h - \ell}{v_b(A'_b)} &\leq \frac{h - \ell}{v_c(A_c^*)} \\ v_b(A'_b) &\geq v_c(A_c^*). \end{aligned}$$

As a result,  $\mathbb{P}(i, A^*) = v_i/v_c(A_c^*) \geq v_i/v_b(A'_b) = \mathbb{P}(i, A')$ , contradicting the fact that  $i$  strictly benefited from participating in the blocking pair.

We conclude there exists no blocking pair  $(b, i)$ , and thus that  $A^*$  is stable.

*SEF1*: Suppose, for contradiction, that  $A^*$  is not SEF1 with respect to the virtual values. This means there exists buyers  $b, c$  so that  $b$  envies  $c$  after every pairwise swap of items, i.e.

$$v_b(A_b^*) + v_{bi} - v_{bj} < v_b(A_c^*) - v_{bi} + v_{bj} \text{ for all } j \in A_b^*, i \in A_c^*.$$

Consider arbitrary  $j \in A_b^*$  and  $i \in A_c^*$  so that  $v_{bj} = \ell$  and  $v_{bi} = h$  (such items exist since  $b$  envies  $c$ ). If it were the case that  $v_{ci} = \ell$ , then it would be possible to increase the product of utilities by exchanging  $i$  and  $j$  ( $b$  is better off and  $c$  is no worse off). The maximality of  $A^*$  thus implies that  $v_{ci} = h$ . We similarly conclude that  $v_{cj} = \ell$ . This means both buyers agree that every item  $j \in A_b^*$  with  $v_{bj} = \ell$  is a low value item, and similarly every  $i \in A_c^*$  with  $v_{bi} = h$  is a high value item.

Let  $H_b = \{k \in \mathcal{I} : v_{bk} = h\}$  and  $L_b = \{k \in \mathcal{I} : v_{bk} = \ell\}$ , and define  $H_c, L_c$  analogously. We've established that  $H_b \cap A_c^* \subseteq H_c \cap A_c^*$  and  $L_b \cap A_b^* \subseteq L_c \cap A_b^*$ .

Construct  $A'$  from  $A^*$  by exchanging  $i$  and  $j$ . So  $A'_b = A_b^* \cup i \setminus j$ ,  $A'_c = A_c^* \cup j \setminus i$  and  $A'_d = A_d^*$  for all  $d \in \mathcal{B} \setminus \{b, c\}$ . Let  $\delta = h - \ell$ .

Since  $A^*$  is not SEF1,  $v_b(A'_b) < v_b(A'_c)$ . Buyer  $c$  receives an  $\ell$  item in exchange for an  $h$  one, so  $v_c(A'_c) = v_c(A_c^*) - \delta < v_c(A_c^*)$ , similarly  $v_b(A'_b) = v_b(A_b^*) + \delta > v_b(A_b^*)$ . Because  $c$  also has high value for those items in  $A'_c$  that  $b$  has high value for,

$$\begin{aligned} v_b(A'_c) &= v_b(A'_c \cap H_b) + v_b(A'_c \cap L_b) \\ &= h \cdot |A'_c \cap H_b| + \ell \cdot |A'_c \cap L_b| \\ &= v_c(A'_c \cap H_b) + \ell \cdot |A'_c \cap L_b| \\ &\leq v_c(A'_c \cap H_b) + v_c(A'_c \cap L_b) = v_c(A'_c). \end{aligned}$$

Putting it all together shows

$$v_b(A_b^*) < v_b(A'_b) < v_b(A'_c) \leq v_c(A'_c) < v_c(A_c^*). \quad (3)$$

Finally,

$$\begin{aligned} v_b(A'_b) \cdot v_c(A'_c) &= (v_b(A_b^*) + \delta)(v_c(A_c^*) - \delta) = v_b(A_b^*) \cdot v_c(A_c^*) + \delta(v_c(A_c^*) - \delta - v_b(A_b^*)) \\ &= v_b(A_b^*) \cdot v_c(A_c^*) + \delta(v_c(A'_c) - v_b(A_b^*)) \\ &> v_b(A_b^*) \cdot v_c(A_c^*), \end{aligned}$$

where the final transition is from Equation (3). This contradicts the fact that  $A^*$  maximized the product of virtual values.

We conclude that  $A^*$  is SEF1 in terms of virtual values. SEF1 in terms of utility follows from the fact that a buyer's utility is monotone in their virtual values.

*Pareto optimality:* This follows directly from  $A^*$  being welfare maximizing.  $\square$

As was the case for identical buyers, not all these properties generalize to arbitrary exposure constraints. Carefully chosen exposure constraints can force a situation where the choice sets of two buyers have to overlap. Suppose, when this happens, buyers disagree on the value of the commonly recommended items and agree on which of the other items have low and high values. Now maximizing buyer welfare will result in allocating the items values highly by both to whoever values the common items least, and the low value items to the other. This may create more envy than what can be eliminated by a single exchange of items. The following example illustrates this.

**EXAMPLE 4.** Consider an instance with  $n = 2, k = 15$  and 20 unique items of three types. There are ten unique items of type a, each with exposure limit two and  $v_{1a} = 2, v_{2a} = 1$ . There are five unique items of type b, each with exposure limit one and  $v_{1b} = 2, v_{2b} = 2$ . There are five items of type c, each with exposure limit one, which both buyers value at one. Table 5 summarizes the instance.

The unique recommendation which maximizes the product of buyers' virtual values (equivalently, expected welfare) is boxed in Table 5: Buyer 1 is recommended all items of type a and c, and buyer 2 all items of types a and b. Buyer 1 has virtual value 26 for her own choice set and 31 for buyer 2's set, so buyer 1 envies buyer 2. Since virtual values are at most two, no single exchange of items can eliminate this envy and the recommendation is not SEF1 with respect to virtual values. By monotonicity of the logarithm, the swap envy will remain when converting the virtual values into expected utilities.

**Table 5** Buyers' virtual values.

Item type	$a$	$b$	$c$	$\omega$
# of items	10	5	5	–
Exposure limit	2	1	1	–
Buyer 1	$\boxed{2}$	2	$\boxed{1}$	$\boxed{1}$
Buyer 2	$\boxed{1}$	$\boxed{2}$	1	$\boxed{1}$

Stability and Pareto optimality, however, can still be guaranteed. The proof that stability holds is largely similar to the unit exposure case, where the existence of a blocking pair contradicted  $A^*$  being welfare maximizing. The main wrinkle is that the seller ejected from the deviating buyer's choice set may already be recommended to the buyer who suffered from the deviation, so care must be taken when constructing the alternative recommendation  $A'$ .

**THEOREM 7.**  *$A^*$  is stable and PO for buyers with dichotomous values.*

## 5. Alternative Algorithms for Constrained Recommendation

Though it is not possible to guarantee stability for general instances, one may attempt to find a stable recommendation on those instances that permit it. In Section EC.6 we present a polynomially-sized integer program which finds stable recommendations when they exist and otherwise minimizes the largest incentive any seller has to deviate. Unfortunately, this approach does not appear to scale to realistic instance sizes.

In this section we propose alternative algorithms accommodating constraints on the number of exposures per product as benchmark against welfare maximizing recommendations. We describe the algorithms under unit constraints, but all are simple to modify to arbitrary exposure constraints.

### 5.1. Round Robin Recommendations

The round robin algorithm fixes a permutation of buyers, then cycles through the buyers  $k$  times. At each step, the active buyer is assigned her highest value item from the remaining items. The only modification required under general exposure constraints is that a buyer may not be assigned the same item twice.

In the instance of Example 1, round robin results in each buyer being recommended one high value item and one low value item. An advantage of round robin is that common high value items are shared among buyers. This increases the probability of high value items being selected by reducing the relative competition they face. This, in turn, reduces their incentive to deviate and promotes stability. We show that round robin, which is known to be EF1 in traditional indivisible goods settings, is also SEF1.

**THEOREM 8.** *Round robin is SEF1.*

Proof. Let  $A$  be the round robin allocation. Suppose for contradiction  $A$  is not SEF1. Then there exists buyers  $b, c \in \mathcal{B}$  so that  $b$  envies  $c$  after any one exchange of items between their allocations. Assume without loss of generality that  $b$  was after  $c$  in the permutation of buyers, otherwise  $b$  would not envy  $c$ . We ignore other buyers, since their bundles do not affect  $b$ 's envy towards  $c$ .

Label  $A_b = \{b_1, b_2, \dots, b_k\}$  and  $A_c = \{c_1, c_2, \dots, c_k\}$ , where items are indexed in the order they are assigned. This implies,  $v_b(b_i) > v_b(b_j)$  and similarly  $v_c(c_i) > v_c(c_j)$  whenever  $j > i$ .

By the nature of round robin, when  $b$  was allocated  $b_i$ , all  $c_j, j > i$  were still unassigned. This implies  $v_b(b_i) \geq v_b(c_{i+1})$  for all  $i \in [k-1]$ . By assumption,  $b$  envies  $c$ , so  $\sum_{i \in [k]} v_b(b_i) < \sum_{i \in [k]} v_b(c_i)$ . Since  $\sum_{i \in [k-1]} v_b(b_i) \geq \sum_{i \in [k-1]} v_b(c_{i+1})$ , it follows that  $v_b(c_1) > v_b(b_k)$ .

Create a new allocation  $A'$  by swapping  $b_k$  and  $c_1$  and leaving the other buyers unchanged, so  $A'_b = A_b \setminus b_k \cup c_1$ ,  $A'_c = A_c \cup b_k \setminus c_1$  and  $A'_d = A_d$  for all  $d \in \mathcal{B} \setminus \{b, c\}$ . Now

$$v_b(A'_b) = v_b(c_1) + \sum_{i \in [k-1]} v_b(b_i) > v_b(b_k) + \sum_{i \in [k-1]} v_b(c_{i+1}) = v_b(A'_c),$$

contradicting that  $b$  has swap envy towards  $c$ . We conclude that round robin is SEF1.  $\square$

Because round robin leads to approximately equal welfare across buyers, choice sets are similarly competitive (or impartial, in the parlance of Theorem 2). This should reduce sellers' incentive to deviate, making stability more likely. Despite being SEF1, round robin need not be Pareto optimal or stable (see Section EC.7).

## 5.2. Online Greedy Top- $k$

This strategy iterates over buyers in a random order and recommends each buyer her top- $k$  (remaining) items. This mimics a natural practice: buyers arrive on a website one at a time, and the platform recommends the set of items that would provide the buyer the highest utility without considering future arrivals. Because greedy top- $k$  considers only a single buyer at a time, its recommendations are based on much less information than those of the two preceding algorithms, both of which consider all buyers and their preferences when constructing a recommendation profile.

Despite the popularity of this approach, it is particularly susceptible to instability, especially when buyers agree that some items are superior to others. Consider the instance in Example 1. The greedy strategy will assign both high value items to first buyer, leaving the other buyer with both low value items. Sellers of both the high value items would get a higher purchase probability by displacing one of the two low value items in the second choice set. The second buyer will accept such displacement because it increases her welfare. As a result, the greedy recommendation is unstable. Identical values are not required for instability: any overlap in buyers' top- $k$  items suffices.

Greedy top- $k$  also fails to satisfy basic fairness and efficiency properties. Consider an instance with two buyers and  $k \geq 3$  and assume that buyer 1 has value 1 for all items. Without loss of generality, suppose that items  $[k]$  are included in buyer 1's choice set. Set buyer 2's values as follows: they value items  $[k]$  highly and the remainder at some low value. Buyer 1 is considered first, without regard for other buyers, after which buyer 2 is recommended items  $\{k + 1, \dots, 2k\}$ , all of which they have low value for. This recommendation is neither Pareto optimal nor SEF1.

PROPOSITION 2. *Greedy top- $k$  recommendations need not be stable, SEF1 or Pareto optimal.*

### 5.3. Online Variants of Batch Algorithms

Greedy top- $k$  is better suited to certain online settings since it only requires a single buyer's preferences to make recommendations, while the preceding algorithms are more suited to batch recommendations. We outline a simple framework for converting a batch recommendation algorithm to an online algorithm by fixing past recommendations and simulating future arrivals using the empirical distributions. The goal is to permit recommending to buyers one at a time as they arrive, while using some population level information to overcome the drawbacks of greedy top- $k$ .

Let  $b_t$  denote the buyer arriving in time step  $t$  as well as her item preferences. By now,  $t - 1$  buyers have received recommendations, let  $\mathcal{B}_{t-1} = \{b_s : s \leq t - 1\}$  be those buyers and  $A_1, \dots, A_{t-1}$  their  $k$ -recommendations. Construct an instance  $\mathcal{I}_t = (\hat{B}_t^n, c^t)$ , where  $\hat{B}_t^n$  consists of  $b_t$  together with  $n - t - 1$  buyers drawn uniformly at random with replacement from  $\mathcal{B}_t$ , and remaining capacities  $c_i^t = c_i - \sum_{j \in [t-1]} \mathbf{1}[i \in A_j]$ . Notice  $c^t$  is exactly the capacities remaining after the first  $t - 1$  recommendations (the past is immutable), and  $\hat{B}_t^n$  assumes future arrivals will have the same distribution as historical ones. If  $\hat{A}_t, \dots, \hat{A}_n$  are the batch recommendations on  $\mathcal{I}_t$ , recommend  $\hat{A}_t$  to buyer  $b_t$ .

For illustration, we evaluate this online variant of round robin. Of course, the online variant of an algorithm provides no more guarantees than the batch version.

## 6. Measuring Instability in Real World Datasets

Theorem 1 shows that there are instances where stability is impossible, but we may hope that such instances are rare. Moreover, even in the absence of stability it is possible that very few sellers participate in blocking pairs, or that they have so little to gain from deviating from the match that it is practically irrelevant. In this section we investigate whether common recommendation strategies lead to stable matches in real-world datasets and, if not, whether sellers have a significant incentive to deviate.

## 6.1. Datasets

We use three datasets for our experiments: two with customers’ ratings on products from Amazon (in the Automotive and Musical Instruments categories) and one with renters’ ratings on clothing from Rent-the-runway (Table 6). These are classic physical goods markets with a natural capacity constraint dictated by inventory levels. Amazon is a platform that matches buyers to third-party sellers. Rent-the-runway is becoming a platform where multiple suppliers maintain their portfolio of garments (Chang 2018). As such, suppliers’ incentive to participate in any recommender system used by these platforms become salient.

**Table 6** Summary of datasets used in the experiments.

Dataset	# users	# items	# ratings	Rating summary			
				Min	Median	Avg	Max
Amazon Automotive	193651	79437	1711519	1	5	4.46	5
Amazon Musical Instrument	27530	10620	231392	1	5	4.47	5
Rent the Runway	105508	5850	192462	2	10	9.09	10

Note: All data are available at <https://cseweb.ucsd.edu/~jmcauley/datasets.html>. We use the full Rent-the-runway dataset. For Amazon, we use the small, dense subsets with buyers and items that occur at least five times.

## 6.2. Experiments

We simulate a platform’s recommendations to a subset of buyers. We take that the platform has estimated rating of each item for each buyer from her past ratings using a learning algorithm (e.g., a collaborative filter).<sup>8</sup> The platform recommends  $k = 5$  items to each buyer using one of the strategies outlined in Section 4 and Section 5. For simplicity, we assume that the platform has only one copy of each item and does not recommend any item to more than one buyer to avoid the risk of stockout.

For the learning algorithm, one can use a content based filter (leveraging both the user and item attributes) or a collaborative filter (using only ratings). Content based filters are less susceptible to cold-start problems, but require rich metadata that may limit their applicability. Collaborative filters require more observations to perform well, but can learn preference on aspects of items that are hard to attribute (e.g., style). It has been shown that collaborative filters with minimal observations per item can perform better than content based filters using metadata (Pilászy and Tikk 2009). We train a matrix factorization based collaborative filters on entire datasets (Zhou

<sup>8</sup> Ratings measure how satisfied a customer has been with a product and are taken to represent a customer’s utility from a product for designing recommender systems (Adomavicius and Tuzhilin 2005).

et al. 2008).<sup>9</sup> Then we randomly select  $B \in \{50, 100, 200\}$  buyers and  $k \times B$  items to form a pool of buyers and items to be matched. The collaborative filter’s predicted ratings for each buyer-item pair is used as the buyer’s expected utility for the item. We take  $u_\omega$  to be 0.<sup>10</sup> We track five metrics for each recommendation strategy:

1. **Welfare:** the average buyer welfare from their recommended choice sets. All else being equal, a platform would like to provide its buyers higher welfare.
2. **Move:** the percentage of sellers who participate in a blocking pair, i.e. can and prefer to form different matches with the buyers. This measures how *widespread* the incentive to deviate is.
3. **Gain:** the average percentage improvement in purchase probability of such sellers if they deviated to maximize their purchase probability. This measures how *strong* the incentive is.
4. **Envy:** the percentage of buyers who envy another. Fewer buyers with envy signals that the recommendations are more fair.
5. **Swap Envy:** the percentage of buyers who envy another *even after their most preferred exchange of items*. Low swap envy implies that whatever envy exists is limited and can be removed by a single swap.

We repeat experiments with 16 random instances from each dataset and report average metrics in Table 7.

There is widespread and substantial incentive to deviate from the recommendations of the greedy top- $k$  algorithm. Nearly all sellers participate in blocking pairs, and they often stand to improve their expected sales by more than 100% by deviating from the system’s recommendations.<sup>11</sup> While most of the sellers would still like to deviate under the other strategies, their potential gain from doing so is substantially lower. Maximizing total buyer welfare appears to be the most stable: roughly half the sellers have incentive to deviate but they stand to gain only about 10-12%. Round robin performs slightly worse, but the simplicity of the algorithm may make it an attractive option when it is not computationally feasible to maximize utility on large datasets. The online variant of round robin compares poorly to the two batch strategies, but is comfortably the best of the online algorithms with the average gain from deviating roughly half that of greedy top- $k$ . Note that it achieves this with the same incremental access to data as the greedy top- $k$ .

<sup>9</sup> We do not set aside a test set for the main experiment since out-of-sample prediction is not the goal. However, a separate evaluation using 20% data for testing shows a RMSE of 1.1 on a 10 point scale for Rent-the-runway and 0.64 – 0.7 on a 5 point scale on the two Amazon datasets. These suggest a reasonably accurate recommender system.

<sup>10</sup> So  $v_\omega = 1$ . The results were quantitatively similar and qualitatively identical without an outside option, i.e.  $v_\omega = 0$ . Setting  $u_\omega$  to average rating of the items in the dataset also leads to qualitatively similar results.

<sup>11</sup> Some caution is warranted in interpreting these gains. We report the average incentive an *individual* seller has to deviate from the current matching. This is not necessarily the gains the sellers will realize if *all of them* deviate to maximize their expected sales. We expect that gains will be lower under this setting, since particularly worse off buyers will attract multiple new sellers, thereby leading to increased competition.



**Table 7** Welfare, instability and envy in 5-recommendations under unit exposure constraints.

	Number of users/items									
	50/250					200/1000				
Automotive	Welfare <sup>+</sup>	Move <sup>-</sup>	Gain <sup>-</sup>	Envy <sup>-</sup>	Swap Envy <sup>-</sup>	Welfare <sup>+</sup>	Move <sup>-</sup>	Gain <sup>-</sup>	Envy <sup>-</sup>	Swap Envy <sup>-</sup>
Max Welfare	6.39 (0.01)	50.1% (1.15%)	8.4% (0.40%)	48.5% (2.67%)	0.4% (0.26%)	6.47 (0.01)	56.6% (0.67%)	8.9% (0.31%)	44.2% (1.75%)	0.4% (0.21%)
Round-robin	6.36 (0.01)	64.8% (1.14%)	11.5% (0.24%)	40.9% (1.69%)	0.0% (0.00%)	6.44 (0.01)	70.8% (0.39%)	11.7% (0.13%)	38.0% (0.65%)	0.0% (0.00%)
Greedy top-5	6.27 (0.01)	96.8% (0.23%)	113.6% (5.40%)	76.0% (1.36%)	61.1% (1.36%)	6.36 (0.01)	99.2% (0.05%)	145.4% (4.49%)	71.7% (0.75%)	58.8% (0.99%)
Online RR	6.27 (0.01)	92.8% (0.53%)	50.9% (3.47%)	64.4% (1.90%)	1.9% (0.37%)	6.36 (0.01)	97.7% (0.26%)	70.4% (4.12%)	51.8% (0.66%)	1.8% (0.17%)
Musical Instr.										
Max Welfare	6.37 (0.01)	48.7% (1.64%)	7.8% (0.37%)	46.8% (2.33%)	0.1% (0.12%)	6.43 (0.01)	55.2% (0.72%)	8.4% (0.27%)	50.1% (1.81%)	0.8% (0.18%)
Round-robin	6.34 (0.01)	64.6% (1.00%)	10.3% (0.26%)	38.7% (1.79%)	0.0% (0.00%)	6.40 (0.01)	70.5% (0.34%)	10.8% (0.14%)	37.4% (0.61%)	0.0% (0.00%)
Greedy top-5	6.26 (0.01)	97.0% (0.18%)	96.0% (5.48%)	75.4% (1.45%)	60.0% (2.33%)	6.33 (0.01)	99.1% (0.07%)	128.1% (5.31%)	73.8% (0.64%)	59.3% (1.01%)
Online RR	6.26 (0.01)	93.6% (0.49%)	44.9% (1.88%)	63.2% (2.18%)	2.1% (0.37%)	6.32 (0.00)	97.7% (0.20%)	65.3% (3.33%)	50.2% (0.97%)	1.4% (0.12%)
Rent-the-runway										
Max Welfare	11.03 (0.01)	55.6% (2.80%)	10.9% (0.98%)	50.7% (3.37%)	0.2% (0.17%)	11.11 (0.00)	63.9% (1.81%)	13.5% (0.92%)	52.7% (1.31%)	0.3% (0.14%)
Round-robin	10.99 (0.01)	70.8% (0.95%)	14.0% (0.23%)	55.6% (2.22%)	0.0% (0.00%)	11.07 (0.00)	73.3% (0.47%)	14.4% (0.13%)	51.3% (1.25%)	0.0% (0.00%)
Greedy top-5	10.88 (0.01)	97.6% (0.12%)	143.0% (4.77%)	87.9% (1.04%)	79.4% (1.48%)	10.96 (0.00)	99.2% (0.06%)	169.0% (4.72%)	87.9% (0.36%)	77.1% (0.57%)
Online RR	10.91 (0.01)	94.3% (0.50%)	52.8% (2.75%)	68.7% (2.21%)	1.1% (0.30%)	10.99 (0.00)	98.1% (0.10%)	84.5% (1.81%)	59.1% (1.49%)	1.6% (0.12%)

Note: 1) Welfare: The average welfare of the buyers from their choice sets. 2) Move: The % of all the sellers who can and would like to move to a different buyer. 3) Gain: The % they would gain in probability of purchase by doing so (over a baseline average of  $\approx 0.2$ ). 4) Envy: The % of buyers who envy the allocation of another. 5) Swap Envy: The % of buyers who envy another even after their most preferred swap. Reported numbers are averages over 16 random draws. The numbers in parenthesis are standard errors. The experiments with 100 users and 500 items produce qualitatively similar results and are omitted. <sup>-</sup>: smaller the better; <sup>+</sup>: larger the better.

Regarding envy, we see that 70-90% of the buyers have envy under the greedy strategy. Moreover, the fraction of the envious buyers does not reduce significantly even when they are allowed to swap an item with the envied buyers. Round robin (both versions) and max welfare do much better. Though roughly 50% of buyers still have envy, the recommendations are nearly SEF1: whatever envy exists can be eliminated by exchanging one pair of items between choice sets.

Unsurprisingly, maximizing total welfare leads to the highest buyer utility. Round robin yields utilities within 1% of optimal. The online algorithms, which only see a single buyer at a time, perform 1-2% worse. Greedy top- $k$ 's myopic decisions are punished by the decreasing marginal utility to adding items to the choice set.

In all, maximizing buyer welfare stands out as with the highest welfare, smallest incentives to deviate from the recommendations and near-zero swap envy.

## 7. Discussion

Personalized recommendations play an important role in matching buyers to sellers in large marketplaces. We initiate a study of the stability of matches formed by such recommendations under a discrete choice model. Though stable matches may not always be possible, we find that maximizing buyer welfare is not only provably stable in restricted preference domains but also leads to relatively low incentives to deviate in real data sets. In this section, we discuss issues not raised elsewhere and avenues for future work.

Many offline matching markets have unraveled in the absence of stability (school-choice, medical labor market, etc.). Due to easier search and discovery, online markets are even more susceptible. Online markets may unravel differently than offline markets. Instead of buyers reaching out to individual sellers, we may see disgruntled sellers multi-homing on various platforms and unsatisfied buyers browsing other platforms before purchasing. These still drive transactions off-platform, which is what platform owners seek to prevent.

Our findings suggest that maximizing buyer welfare is reasonable when the objective is a stable marketplace. But stability of one marketplace may not be enough in an ecosystem where multiple marketplaces/platforms compete. There is a growing interest in fairness towards certain groups at a marketplace (e.g., exposure for minority owned or small businesses). An explicit consideration of buyer, seller, and platform objectives (emphasized in recent recommender systems literature) may be necessary to make a platform successful in this environment.

We have not considered strategy-proof mechanisms in this study. Typically, recommender systems estimate preferences from buyers' past behavior. These values are somewhat resistant to manipulation, since manipulation would require buyers modifying their browsing, search or purchases for some time. In an alternative setting where agents directly report preferences primarily for matching, questions of strategyproofness and incentive compatibility become salient.

Several interesting questions remain open: Are there scalable methods that find stable recommendations when they exist and otherwise minimize the incentive to deviate from the recommendations? The integer program in Section EC.6 is a step in this direction but scalability remains an issue. We ignored issues of competition but it may be important to understand, for example, how a platform's recommendation algorithm should change when a rival actively attempts to lure away users. Finally, is there a complete characterization of exactly which instances or preference structures permit stable recommendations? Do real-world instances satisfy these conditions?

The stability of personalized recommendations is a topic of enormous practical relevance and theoretical significance. We hope this study leads to future research on this important topic.

## References

- H. Abdollahpouri, G. Adomavicius, R. Burke, I. Guy, D. Jannach, T. Kamishima, J. Krasnodebski, and L. Pizzato. 2020. Multistakeholder recommendation: Survey and research directions. *User Modeling and User-Adapted Interaction* 30, 1 (2020), 127–158.
- H. Abdollahpouri and R. Burke. 2019. Multi-stakeholder recommendation and its connection to multi-sided fairness. *arXiv preprint arXiv:1907.13158* (2019).
- A. Abdulkadiroglu and T. Sönmez. 2013. Matching markets: Theory and practice. *Advances in Economics and Econometrics* 1 (2013), 3–47.
- G. Adomavicius and A. Tuzhilin. 2005. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE transactions on knowledge and data engineering* 17, 6 (2005), 734–749.
- J. W. Alba and J. W. Hutchinson. 1987. Dimensions of consumer expertise. *Journal of Consumer Research* 13, 4 (1987), 411–454.
- A. Alkan, G. Demange, and D. Gale. 1991. Fair allocation of indivisible goods and criteria of justice. *Econometrica* 59, 4 (1991), 1023–1039.
- Y. Bakos and H. Halaburda. 2020. Platform competition with ultrahoming on both Sides: subsidize or not? *Management Science* 66, 12 (2020), 5599–5607.
- K. Bando. 2012. Many-to-one matching markets with externalities among firms. *Journal of Mathematical Economics* 48, 1 (2012), 14–20.
- K. Bando. 2014. A modified deferred acceptance algorithm for many-to-one matching markets with externalities among firms. *Journal of Mathematical Economics* 52 (2014), 173–181.
- K. Bando, R. Kawasaki, and S. Muto. 2016. Two-sided matching with externalities: A survey. *Journal of the Operations Research Society of Japan* 59, 1 (2016), 35–71.
- M. Bateni, Y. Chen, D. F. Ciocan, and V. Mirrokni. 2022. Fair resource allocation in a volatile marketplace. *Operations Research* 70, 1 (2022), 288–308.
- O. Ben-Porat and M. Tennenholtz. 2018. A game-theoretic approach to recommendation systems with strategic content providers. *Proceedings of the 32nd Annual Conference on Neural Information Processing Systems (NeurIPS)* 31 (2018).
- G. Benadè, A. Kazachkov, C.-A. Psomas, and A. D. Procaccia. 2018. How to Make Envy Vanish Over Time. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*. 593–610.
- G. Benadè, A. Kazachkov, C.-A. Psomas, A. D. Procaccia, and D. Zeng. 2022. Fair and Efficient Online Allocations. Manuscript.
- S. J. Brams and A. D. Taylor. 1995. An Envy-Free Cake Division Protocol. *The American Mathematical Monthly* 102, 1 (1995), 9–18.

- K. A. Bryan and J. S. Gans. 2019. A theory of multihoming in rideshare competition. *Journal of Economics & Management Strategy* 28, 1 (2019), 89–96.
- R. D. Burke, H. Abdollahpouri, B. Mobasher, and T. Gupta. 2016. Towards Multi-Stakeholder Utility Evaluation of Recommender Systems. *UMAP (Extended Proceedings)* 750 (2016).
- I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)* 7, 3 (2019), 1–32.
- I. Caragiannis and S. Narang. 2022. Repeatedly Matching Items to Agents Fairly and Efficiently. *arXiv preprint arXiv:2207.01589* (2022).
- M. Carroll, D. Hadfield-Menell, S. Russell, and A. Dragan. 2021. Estimating and penalizing preference shift in recommender systems. In *Proceedings of the 15th ACM Conference on Recommender Systems*. 661–667.
- A. Chakraborty, A. Hannak, A. J. Biega, and K. Gummadi. 2017. Fair sharing for sharing economy platforms. In *Proceedings of FATREC Workshop on Responsible Recommendation at RecSys 2017*.
- L. Chang. 2018. Rent the Runway is launching a new way for designers to partner with the company — and it’s great news for subscribers.
- X. Chen, S. Li, H. Li, S. Jiang, Y. Qi, and L. Song. 2019. Generative adversarial user model for reinforcement learning based recommendation system. In *Proceedings of the 16th International Conference on Machine Learning (ICML)*. 1052–1061.
- G. De Jong, A. Daly, M. Pieters, and T. Van der Hoorn. 2007. The logsum as an evaluation measure: Review of the literature and new results. *Transportation Research Part A: Policy and Practice* 41, 9 (2007), 874–889.
- B. Dutta and J. Massó. 1997. Stability of matchings when individuals have preferences over colleagues. *journal of economic theory* 75, 2 (1997), 464–475.
- F. Echenique and J. Oviedo. 2004. Core many-to-one matchings by fixed-point Methods. *Journal of Economic Theory* 115, 2 (2004), 358–376.
- F. Echenique and J. Oviedo. 2006. A theory of stability in many-to-many matching markets. *Theoretical Economics* 1 (2006), 233–273.
- F. Echenique and M. B. Yenmez. 2007. A solution to matching with preferences over colleagues. *Games and Economic Behavior* 59, 1 (2007), 46–71.
- Edmund Eisenberg and David Gale. 1959. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics* 30, 1 (1959), 165–168.
- M. D. Ekstrand, A. Das, R. Burke, and F. Diaz. 2022. Fairness in Recommender Systems. In *Recommender Systems Handbook*. Springer-Verlag, 679–707.

- F. Eskandarian and B. Mobasher. 2020. Using stable matching to optimize the balance between accuracy and diversity in recommendation. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization*. 71–79.
- D. Fleder and K. Hosanagar. 2009. Blockbuster culture’s next rise or fall: The impact of recommender systems on sales diversity. *Management science* 55, 5 (2009), 697–712.
- D. M. Fleder and K. Hosanagar. 2007. Recommender systems and their impact on sales diversity. In *Proceedings of the 8th ACM Conference on Economics and Computation (EC)*. 192–199.
- D. Foley. 1967. Resource allocation and the public sector. *Yale Economics Essays* 7 (1967), 45–98.
- R. Freeman, E. Micha, and N. Shah. 2021. Two-sided Matching Meets Fair Division. *arXiv preprint arXiv:2107.07404* (2021).
- D. Gale and L. S. Shapley. 1962. College Admissions and the Stability of Marriage. *American Mathematical Monthly* 69, 1 (1962), 9–15.
- S. Gollapudi, K. Kollias, and B. Plaut. 2020. Almost Envy-Free Repeated Matching in Two-sided Markets. In *International Conference on Web and Internet Economics*. 3–16.
- I. Goodfellow, I. Bengio, and I. Courville. 2016. *Deep learning*. MIT Press.
- J. R. Hauser and B. Wernerfelt. 1990. An evaluation cost model of consideration sets. *Journal of consumer research* 16, 4 (1990), 393–408.
- Z. Huang, M. Li, X. Shu, and T. Wei. 2022. Online Nash Welfare Maximization Without Predictions. *arXiv preprint arXiv:2211.03077*. (2022).
- A. Igarashi, Y. Kawase, W. Suksompong, and H. Sumita. 2022. Fair Division with Two-Sided Preferences. *arXiv preprint arXiv:2206.05879* (2022).
- A. S. Kelso and V. P. Crawford. 1982. Job Matching, Coalition Formation, and Gross Substitutes. *Econometrica: Journal of the Econometric Society* (1982), 1483–1504.
- R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. 2004. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 6th ACM Conference on Economics and Computation (EC)*. 125–131.
- R. Makhijani, S. Chakrabarti, D. Struble, and Y. Liu. 2019. LORE: A large-scale offer recommendation engine with eligibility and capacity constraints. In *Proceedings of the 13th ACM Conference on Recommender Systems*. 160–168.
- C. F. Manski. 1977. The structure of random utility models. *Theory and decision* 8, 3 (1977), 229.
- D. McFadden. 1973. Conditional logit analysis of qualitative choice behavior. *Frontiers in Econometrics* (1973).
- T. Moins, D. Aloise, and S. J. Blanchard. 2020. RecSeats: A hybrid convolutional neural network choice model for seat recommendations at reserved seating venues. In *Proceedings of the 14th ACM Conference on Recommender Systems*. 309–317.

- A. Mumcu and I. Saglam. 2010. Stable one-to-one matchings with externalities. *Mathematical Social Sciences* 60, 2 (2010), 154–159.
- J. Nash. 1950. The bargaining problem. *Econometrica* 18, 2 (1950), 155–162.
- P. Nguyen, J. Dines, and J. Krasnodebski. 2017. A Multi-Objective Learning to Re-rank Approach to Optimize Online Marketplaces for Multiple Stakeholders. *arXiv preprint arXiv:1708.00651* (2017).
- G. K. Patro, A. Biswas, N. Ganguly, K. P. Gummadi, and A. Chakraborty. 2020. Fairrec: Two-sided Fairness for Personalized Recommendations in Two-sided Platforms. In *Proceedings of The Web Conference 2020*. 1194–1204.
- I. Pilászy and D. Tikk. 2009. Recommending New Movies: Even a Few Ratings Are More Valuable than Metadata. In *Proceedings of the Third ACM Conference on Recommender Systems*. 93–100.
- A. D. Procaccia. 2016. Cake Cutting Algorithms. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press, Chapter 13.
- M. Pycia. 2012. Stability and preference alignment in matching and coalition formation. *Econometrica* 80, 1 (2012), 323–362.
- M. Pycia and M. B. Yenmez. 2022. Matching with externalities. *The Review of Economic Studies* 90, 2 (06 2022), 948–974.
- A. E. Roth. 1984a. The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of political Economy* 92, 6 (1984), 991–1016.
- A. E. Roth. 1984b. Stability and polarization of interests in job matching. *Econometrica: Journal of the Econometric Society* (1984), 47–57.
- A. E. Roth. 2007. The art of designing markets. *Harvard Business Review* 85, 10 (2007), 118.
- A. E. Roth. 2015. *Who gets what—and why: the new economics of matchmaking and market design*. Houghton Mifflin Harcourt.
- A. E. Roth and M. Sotomayor. 1990. *Two-Sided Matching: A Study in Game-Theoretic Modelling and Analysis*. Cambridge University Press.
- H. Sasaki and M. Toda. 1996. Two-sided matching problems with externalities. *Journal of Economic Theory* 70, 1 (1996), 93–108.
- D. Serbos, S. Qi, N. Mamoulis, E. Pitoura, and P. Tsaparas. 2017. Fairness in package-to-group recommendations. In *Proceedings of the 26th International World Wide Web Conference (WWW)*. 371–379.
- S. M. Shugan. 1980. The cost of thinking. *Journal of Consumer Research* 7, 2 (1980), 99–111.
- Y. Song, Z.n Li, and N. Sahoo. 2022. Matching returning donors to projects on philanthropic crowdfunding platforms. *Management Science* 68, 1 (2022), 355–375.
- Y. Song, N. Sahoo, and E. Ofek. 2019. When and how to diversify—a multicategory utility model for personalized content recommendation. *Management Science* 65, 8 (2019), 3737–3757.

- J. M. Stalnaker. 1953. The Matching Program for Intern Placement. *Journal of Medical Education* 28, 11 (1953), 13–19.
- Ö. Sürer, R. Burke, and E. C. Malthouse. 2018. Multistakeholder recommendation with provider constraints. In *Proceedings of the 12th ACM Conference on Recommender Systems*. 54–62.
- M. Tennenholtz and O. Kurland. 2019. Rethinking search engines and recommendation systems: a game theoretic perspective. *Commun. ACM* 62, 12 (2019), 66–75.
- K. E. Train. 2009. *Discrete choice methods with simulation*. Cambridge University Press.
- H. C. W. L. Williams. 1977. On the formation of travel demand models and economic evaluation measures of user benefit. *Environment and planning A* 9, 3 (1977), 285–344.
- D. Zeng and A. Psomas. 2020. Fairness-efficiency tradeoffs in dynamic fair division. In *Proceedings of the 21st ACM Conference on Economics and Computation*. 911–912.
- Y. Zhou, D. Wilkinson, R. Schreiber, and R. Pan. 2008. Large-scale parallel collaborative filtering for the netflix prize. In *Algorithmic Aspects in Information and Management: 4th International Conference*. 337–348.

## Proofs Omitted From The Main Body

### EC.1. Incentive to Deviate

EXAMPLE EC.1. Consider the instance in Table EC.1 with exposure constraints equal to 1. Note the the instance is 2-balanced. The recommendation  $A_1 = \{c, d\}$  is  $\beta$ -impartial, with  $\beta = \frac{2}{2-\varepsilon}$ , and  $(1, a)$  is a blocking pair, The purchase probability of item  $a$  is  $1/2$  in  $A_1$ . After deviating

**Table EC.1** Buyers' virtual values.

	$a$	$b$	$c$	$d$
Buyer 1	2	1	$2 - \varepsilon$	0
Buyer 2	1	1	1	0

to  $A'_1 = \{a, d\}$ , the purchase probability of  $a$  is 1. The difference between seller  $a$ 's multiplicative improvement and the bound of  $\alpha\beta$  in Theorem 2 is

$$\alpha\beta - \frac{P(a, A')}{P(a, A)} = 2 \cdot \frac{2}{2-\varepsilon} - \frac{1}{1/2} = \frac{4-4+2\varepsilon}{2-\varepsilon} = \frac{2\varepsilon}{2-\varepsilon} \rightarrow 0 \text{ as } \varepsilon \rightarrow 0.$$

By considering the deviation to  $A''_1 = \{a, c\}$ , we similarly find that the upper bound on buyer improvement is tight. Specifically,

$$\alpha\beta - \frac{v_1(A'')}{v_1(A)} = 2 \cdot \frac{2}{2-\varepsilon} - \frac{4-\varepsilon}{2-\varepsilon} = \frac{\varepsilon}{2-\varepsilon} \rightarrow 0 \text{ as } \varepsilon \rightarrow 0.$$

### EC.2. Other restricted preference domains

One may hope to escape this non-existence by placing restrictions on buyer preferences. For two common forms of structured preferences this is unsuccessful. First, suppose each buyer  $b$  is associated with a characteristic vector  $\beta_b$ , each item  $i$  with similar vector  $\gamma_i$ , and  $u_{bi} = \langle \beta_b, \gamma_i \rangle$ . A stable recommendation still need not exist since the instance of Theorem 1 can be factorized. Second, when buyers have identical preference orders over the items (but potentially different values) a stable recommendation always exists under unit exposures for instances with  $n = 2, k = 2$ , but not in general (see Section EC.5.1).

### EC.3. Identical Buyers

THEOREM 5.  $A^*$  is PO and SEF1 for buyers with identical preferences.

Proof. Pareto optimality again follows from maximizing expected buyer welfare.

Suppose for contradiction that  $A^*$  is not SEF1. Then there exists buyers  $b, c$ , so that  $b$  envies  $c$  even any feasible exchange of items between  $A_b^*$  and  $A_c^*$ . Let  $X = A_b^* \cap A_c^*$  denote the items



recommended to both  $b$  and  $c$ . An exchange of items  $i \in X$  and  $j \in A_c^*$  can only be feasible if  $j = i$ , otherwise  $c$  ends up being recommended  $i$  twice. Such an exchange does not change buyer bundles or utility, so we may safely ignore them.

Let  $S = \{(i, j) : i \in A_b^* \setminus X, j \in A_c^* \setminus X, v_i < v_j\}$  denote the set of feasible exchanges that are (strictly) improving for  $b$ . Suppose  $S = \emptyset$ . If  $|X| = k$ , then  $b$  and  $c$  are recommended identical choice sets, and there is no envy. We may conclude that  $|X| < k$ . Since  $k = |A_b^*| = |A_c^*|$ , it follows that  $|A_b^* \setminus X| = |A_c^* \setminus X| > 0$ . Let  $i^- = \arg \min_i \{v_i : i \in A_b^* \setminus X\}$  and  $j^+ = \arg \max_j \{v_j : j \in A_c^* \setminus X\}$ . If  $S = \emptyset$ , then in particular  $(i^-, j^+) \notin S$  and, since this is a feasible exchange, it follows that  $v_{i^-} > v_{j^+}$ . Then

$$\begin{aligned} v(A_b^*) &\geq v(X) + v_{i^-} \cdot |A_b^* \setminus X| > v(X) + v_{j^+} \cdot |A_b^* \setminus X| \\ &= v(X) + v_{j^+} \cdot |A_c^* \setminus X| = v(A_c^*), \end{aligned}$$

contradicting that  $b$  envies  $c$ . It follows that  $S \neq \emptyset$ .

Select arbitrary  $(i, j) \in S$ . By assumption,  $v(A_b^*) + v_j - v_i < v(A_c^*) + v_i - v_j$ . Construct  $A'$  by exchanging  $i$  and  $j$  and keeping the rest of the recommendation unchanged, so  $A'_b = A_b^* \cup \{j\} \setminus \{i\}$  and  $A'_c = A_c^* \cup \{i\} \setminus \{j\}$ .  $A'_d = A_d^*$  for all  $d \in \mathcal{B} \setminus \{b, c\}$ .

Set  $\delta = v_j - v_i > 0$ , so  $v(A'_b) = v(A_b^*) + \delta$  and  $v(A'_c) = v(A_c^*) - \delta$ . Now

$$v(A'_b) \cdot v(A'_c) = (v(A_b^*) + \delta)(v(A_c^*) - \delta) = v(A_b^*) \cdot v(A_c^*) + \delta(v(A_c^*) - (v(A_b^*) + \delta)) \geq v(A_b^*) \cdot v(A_c^*),$$

where the final transition follows from

$$v(A_c^*) - (v(A_b^*) + \delta) \geq v(A'_c) - v(A'_b) > 0$$

since  $A^*$  is not SEF1 by assumption. This contradicts the fact that  $A^*$  maximized the product of virtual values. We conclude that  $A^*$  is SEF1.  $\square$

#### EC.4. Dichotomous buyers

**THEOREM 7.**  *$A^*$  is stable and PO for buyers with dichotomous values.*

*Proof.*

As before, set  $\ell = e^a$  and  $h = e^{a'}$ . We show that maximizing the product of buyer virtual values, which is equivalent to maximizing buyer welfare, is stable (Pareto optimality follows from maximizing buyer welfare).

Assume, for contradiction,  $A^*$  is not stable. Then there exists a blocking pair  $(b, i) \in \mathcal{B} \times \mathcal{I}$ . Let  $c$  be the buyer to whom (the relevant copy of)  $i$  is currently recommended.

Since  $b$  is willing to participate in the blocking pair,  $v_{bi} = h$  and her choice set contains at least one low-valued item. Let  $j \in A_b^*$  be such an item with  $v_{bj} = \ell$ . Construct  $A'$  from  $A^*$  by transferring

$i$  from  $c$  to  $b$ 's choice sets, and completing  $c$ 's choice set by recommending some item  $j'$  that is below capacity after the transfer. Note that  $j'$  must exist, by the definition of a blocking pair, and  $j'$  need not be  $j$ , in particular, when  $j \in A_c^*$  using  $j' = j$  is infeasible. Now  $A'_b = A_b^* + i - j$ ,  $A'_c = A_c^* + j' - i$  and  $A'_d = A_d^*$  for all  $d \in \mathcal{B} \setminus \{b, c\}$ .

We now consider the possible values of  $(v_{bi}, v_{cj'})$ :

1.  $(\ell, \cdot)$ : Now  $v_b(A'_b) \leq v_b(A_b^*)$ . This contradicts  $(b, i)$  being a blocking pair, since  $b$  must strictly gain from participating in a blocking pair and can not do so if  $v_{bi} = \ell = v_{bj}$ .

2.  $(h, h)$ : Now  $v_b(A'_b) > v_b(A_b^*)$  since  $v_{bi} > v_{bj}$  and  $v_c(A'_c) \geq v_c(A_c^*)$  since  $v_{cj'} = h \geq v_{ci}$ . This contradicts that  $A^*$  maximizes the product of virtual values.

3.  $(h, \ell)$ : Now  $v_b(A'_b) = v_b(A_b^*) - \ell + h > v_b(A_b^*)$ . We will handle the cases of  $v_{ci} = \ell$  and  $v_{ci} = h$  separately.

First, suppose  $v_{ci} = \ell$ . Then  $v_c(A'_c) = v_c(A_c^*)$ , implying  $v_c(A'_c) \cdot v_b(A'_b) > v_c(A_c^*) \cdot v_b(A_b^*)$ , contradicting  $A^*$  maximizing the product of virtual values.

Suppose instead  $v_{ci} = h$ . Now  $v_c(A'_c) = v_c(A_c^*) + \ell - h < v_c(A_c^*)$ .

We know that  $v_b(A'_b) < v_b(A_b^*) + h - \ell = v_b(A'_b) < v_c(A_c^*)$ , because both  $b$  and  $i$  benefit from participating in the blocking pair and  $v_{bi} = h = v_{ci}$  by assumption. As a result,  $v_c(A_c^*) - v_b(A_b^*) > h + \ell$ . In contrast,  $v_c(A_c^*) - v_c(A'_c) = h + \ell$ . It follows that  $v_b(A_b^*) < v_c(A'_c)$ , and the resulting product of virtual values of

$$\begin{aligned} v_b(A'_b) \cdot v_c(A'_c) &= (v_b(A_b^*) + h - \ell)(v_c(A_c^*) - h + \ell) \\ &= v_b(A_b^*) \cdot v_c(A_c^*) + (h - \ell)[v_c(A_c^*) - h + \ell - v_b(A_b^*)] \\ &= v_b(A_b^*) \cdot v_c(A_c^*) + (h - \ell)[v_c(A'_c) - v_b(A_b^*)] \\ &> v_b(A_b^*) \cdot v_c(A_c^*) \end{aligned}$$

contradicts the fact that  $A^*$  was welfare maximizing.

We conclude there exists no blocking pair  $(b, i)$ . Hence,  $A^*$  is stable.  $\square$

## EC.5. Other Preference Domains

### EC.5.1. Latent factor models

First, we show that the instance of Theorem 1 (Table 2) can be factorised, which implies that stability can not be guaranteed when buyer values come from a latent factor model.

EXAMPLE EC.2. A stable matching need not exist when values come from a latent factor model. Let  $n = 2, k = 2, f = 2$ . Consider  $\beta_1 = (0.6, 1.4), \beta_2 = (1.7, 0.3)$  and  $\sigma_a = (1.1, 1.2), \sigma_b = (1.3, -0.5), \sigma_c = (0.6, 1.2), \sigma_d = (0.7, 1)$ . The resulting value and utility matrices are shown below.

It is straightforward to verify that this instance allows the same blocking pairs identified in Theorem 1.

**Table EC.2** Utility matrix (left) and valuation matrix (right).

	a	b	c	d		a	b	c	d
1	2.34	0.08	2.04	1.82	1	10.39	1.08	7.69	6.17
2	2.23	2.06	1.38	1.49	2	9.39	7.84	3.97	4.44

### EC.5.2. Identical preference orders

Next, we consider the case where buyers have identical preference orders over the items, but not identical values.

The following example with two buyers, six items, unit exposure constraints and  $k = 3$  shows that stability can not be guaranteed.

**Table EC.3** Buyers' virtual values: identical preference orders do not guarantee stable recommendations existing.

	a	b	c	d	e	f
1	1	4	5	6	7	10
2	0.5	1.7	4.5	5	9	10

However, stability can be guaranteed in the restricted case of unit exposure constraints with  $n = 2 = k$ ,  $m = 4$ .

**PROPOSITION EC.1.** *Consider two buyers with identical preference orders and four items  $\{a, b, c, d\}$  labeled in decreasing order of their virtual values. When  $k = 2$  under unit exposure constraints, at least one of  $\{\{a, d\}, \{b, c\}\}$  and  $\{\{b, c\}, \{a, d\}\}$  is stable.*

*Proof.* Let  $A_1 = \{\{a, d\}, \{b, c\}\}$  and  $A_2 = \{\{b, c\}, \{a, d\}\}$ .

First, we show that in either  $A_1$  or  $A_2$  seller  $a$  is unwilling to participate in a blocking pair. Suppose this is not the case, and  $a$  participates in a blocking pair in both  $A_1$  and  $A_2$ . Then, from  $A_1$ , we conclude  $\frac{v_{1a}}{v_{1a}+v_{1d}} < \frac{v_{2a}}{v_{2a}+v_{2c}}$ . Similarly, from  $A_2$ ,  $\frac{v_{2a}}{v_{2a}+v_{2d}} < \frac{v_{1a}}{v_{1a}+v_{1c}}$ . Together, it follows that

$$\frac{v_{2a}}{v_{2a}+v_{2d}} < \frac{v_{1a}}{v_{1a}+v_{1c}} < \frac{v_{1a}}{v_{1a}+v_{1d}} < \frac{v_{2a}}{v_{2a}+v_{2c}} < \frac{v_{2a}}{v_{2a}+v_{2d}},$$

a contradiction. We conclude  $a$  is unwilling to participate in a blocking pair in at least one of  $A_1, A_2$ ; relabel the buyers so that this happens in  $A_1$ .

Consider  $A_1 = \{\{a, d\}, \{b, c\}\}$ . By construction,  $a$  does not participate in a blocking pair. Item  $b$  does not participate in a blocking pair, since it is currently selected with probability greater than 0.5 but this drops to less than 0.5 after deviating. Item  $d$  does not participate in a blocking pair, since it is the least liked item and can not displace another item from a choice set. Assume that  $c$  participates in a blocking pair, otherwise we have found a stable recommendation. It follows that

$$\frac{v_{2c}}{v_{2b}+v_{2c}} < \frac{v_{1c}}{v_{1a}+v_{1c}}. \tag{EC.1}$$

Now consider  $A_2 = \{\{b, c\}, \{a, d\}\}$ . By the same reasoning as before,  $b$  and  $d$  do not participate in blocking pairs. We investigate whether  $a$  or  $c$  can participate in a blocking pair.

1. Suppose  $a$  participates in a blocking pair. Then  $\frac{v_{2a}}{v_{2a}+v_{2d}} < \frac{v_{1a}}{v_{1a}+v_{1c}}$ . Equivalently,

$$\frac{v_{2d}}{v_{2a}+v_{2d}} > \frac{v_{1c}}{v_{1a}+v_{1c}} \stackrel{(\text{By Eq. EC.1})}{>} \frac{v_{2c}}{v_{2b}+v_{2c}} > \frac{v_{2d}}{v_{2b}+v_{2d}} > \frac{v_{2c}}{v_{2a}+v_{2c}},$$

a contradiction.

2. Suppose  $c$  participates in a blocking pair. Then  $\frac{v_{1c}}{v_{1c}+v_{1b}} < \frac{v_{2c}}{v_{2a}+v_{2c}}$ . It follows that

$$\frac{v_{1c}}{v_{1c}+v_{1b}} < \frac{v_{2c}}{v_{2a}+v_{2c}} < \frac{v_{2c}}{v_{2b}+v_{2c}} \stackrel{(\text{By Eq. EC.1})}{<} \frac{v_{1c}}{v_{1a}+v_{1c}} < \frac{v_{1c}}{v_{1c}+v_{1b}},$$

a contradiction.

We conclude either  $A_1$  is stable or, when  $c$  participates in a blocking pair in  $A_1$ ,  $A_2$  is stable.

□

## EC.6. Finding Stable Recommendations, When They Exist

Example 1 shows that stability can not be guaranteed for general preferences. However, there may still be many instances that permit stable recommendations. We construct an integer program to find a stable match if it exists and, if not, returns recommendations in which the sellers' benefit from participating in a blocking pair is as small as possible.

There are some obstacles to overcome. Stability depends on items' purchase probabilities, which are inherently nonlinear. One option is to assign choice sets to buyers and precompute all the resulting purchase probabilities, however, this leads to an exponentially sized program. We present below a formulation with  $O(|\mathcal{B}|^2 \cdot |\mathcal{I}|^2) = O(|\mathcal{B}|^4 \cdot k^2)$  variables and as many constraints.

Define binary variable  $x_{bi}$  which takes value 1 exactly when  $i \in \mathcal{I}$  is recommended to  $b \in \mathcal{B}$ . The following constraints ensures  $k$  sellers are recommended to each buyer and that the recommendation satisfies capacity constraints

$$\sum_{b \in \mathcal{B}} x_{bi} = c_i, \forall i \in \mathcal{I}, \tag{EC.2}$$

$$\sum_{i \in \mathcal{I}} x_{bi} = k, \forall b \in \mathcal{B}. \tag{EC.3}$$

Define continuous variable  $g \geq 0$  to capture the maximum multiplicative improvement any seller can get from deviating from any solution  $x$ . Consider arbitrary  $(b, i) \in \mathcal{B} \times \mathcal{I}$  and let  $c$  be the buyer that  $i$  is recommended to and  $j$  seller currently recommended to  $b$  who can be displaced by  $i$  (i.e.  $v_{bi} > v_{bj}$ ).

Then

$$g \geq \frac{v_{bi}}{\sum_{l \in \mathcal{I}} x_{bl} v_{bl} - v_{bj} + v_{bi}} \bigg/ \frac{v_{ci}}{\sum_{l \in \mathcal{I}} x_{cl} v_{cl}},$$

where the numerator is the purchase probability of  $i$  after transferring into  $b$ 's bundle and the denominator is their current purchase probability with  $c$ . We rewrite this as

$$g \cdot \left( \sum_{l \in \mathcal{I}} x_{bl} v_{bl} - v_{bj} + v_{bi} \right) \geq \frac{v_{bi}}{v_{ci}} \cdot \left( \sum_{l \in \mathcal{I}} x_{cl} v_{cl} \right).$$

Define continuous variable  $z_{bi} = g \cdot x_{bi}$  for all  $b \in \mathcal{B}, i \in \mathcal{I}$ . To ensure that  $z_{bi}$  takes the appropriate values, we require constraints

$$z_{bi} \geq 0, \tag{EC.4}$$

$$z_{bi} \leq x_{bi} \cdot G, \tag{EC.5}$$

$$z_{bi} \leq g, \tag{EC.6}$$

$$z_{bi} \geq g + (x_{bi} - 1)G, \tag{EC.7}$$

for all  $b \in \mathcal{B}, i \in \mathcal{I}$  and some upper bound  $G$  on  $g$ . Substituting into the above we obtain

$$\sum_{l \in \mathcal{I}} z_{bl} v_{bl} - g v_{bj} + g v_{bi} \geq \frac{v_{bi}}{v_{ci}} \cdot \sum_{l \in \mathcal{I}} x_{cl} v_{cl}, \tag{EC.8}$$

which should hold for all  $b \neq c \in \mathcal{B}, i \neq j \in \mathcal{I}$  as long as  $x_{bj} = 1 = x_{ci}$  and  $v_{bi} > v_{bj}$ . To enforce this we create a new indicator variable

$$\delta_{bj}^{ci} = \begin{cases} 1, & \text{when } x_{bj} = 1 = x_{ci} \text{ and } v_{bi} > v_{bj}, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Now we can rewrite eq. (EC.8) as

$$\sum_{l \in \mathcal{I}} z_{bl} v_{bl} - g v_{bj} + g v_{bi} \geq \frac{v_{bi}}{v_{ci}} \cdot \sum_{l \in \mathcal{I}} x_{cl} v_{cl} - (1 - \delta_{bj}^{ci})M. \tag{EC.9}$$

The following two constraints ensure that  $\delta_{bj}^{ci}$  takes on the value 1 when expected,

$$(1 - \delta_{bj}^{ci})M' \geq v_{bj} - v_{bi} - (x_{bj} + x_{bi} - 2)M, \tag{EC.10}$$

$$-\delta_{bj}^{ci}M' \leq v_{bj} - v_{bi} - (x_{bj} + x_{bi} - 2)M, \tag{EC.11}$$

for  $M' > 2M$ . Observe that when  $x_{bj} = 1 = x_{ci}$  and  $v_{bi} > v_{bj}$ , eq. (EC.10) does not bind and eq. (EC.11) becomes  $-\delta_{bj}^{ci}M' < 0$ , implying  $\delta_{bj}^{ci} = 1$ . When  $x_{ci} + x_{cj} < 2$ , eq. (EC.10) becomes  $(1 - \delta_{bj}^{ci})M' \geq 0$ , ensuring  $\delta_{bj}^{ci} = 0$ , while (EC.11) does not bind. Similarly  $v_{bi} \leq v_{bj}$ , implies  $\delta_{bj}^{ci} = 0$ .

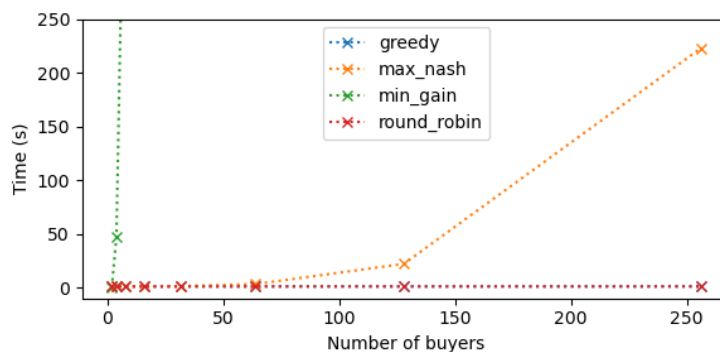
This yields the mixed integer program

$$\begin{aligned} & \min g \\ \text{s.t.} \quad & \text{eqs. (EC.2) to (EC.3)} && \text{(assignment constraints)} \\ & \text{eqs. (EC.4) to (EC.7)} \forall b \in \mathcal{B}, i \in \mathcal{I} && \text{(linearization constraints)} \\ & \text{eqs. (EC.9) to (EC.11)} \forall b \neq c \in \mathcal{B}, i \neq j \in \mathcal{I} && \text{(stability constraints)} \\ & g \geq 0, x \in \{0, 1\}^{|\mathcal{B}| \times |\mathcal{I}|}, z \in \{0, 1\}^{|\mathcal{B}| \times |\mathcal{I}|} \\ & \delta_{bj}^{ci} \in \{0, 1\}^{|\mathcal{B}| \times |\mathcal{I}| \times |\mathcal{B}| \times |\mathcal{I}|} \end{aligned}$$

When  $g \leq 1$ , no seller can improve their purchase probability by participating in a blocking pair and the recommendation is stable.

### EC.6.1. Scalability

We compare the computational cost of the integer program with the three recommendation strategies in Section 5 by training SVD++, a matrix factorization based collaborative filter, on the datasets described in Section 6.1. An instance is created setting  $k = 3$  and randomly selecting  $B \in \{2, 2^2, \dots, 2^8\}$  buyers and a corresponding number of random items and taking the collaborative filter’s estimated buyer-item ratings as values. The time it takes each approach to yield a recommendation is visually represented in Figure EC.1. Clearly, the integer program does not scale to reasonable sizes; we exclude it from further experiments. Maximizing buyer welfare performs well on these tests, but may eventually present computational challenges. Round robin and greedy top- $k$  is consistently extremely fast for all problem sizes.



**Figure EC.1** Time to find recommendations as function of the number of buyers with  $k = 5$ .

### EC.7. Round robin is not stable or PO

Consider the instance with 2 buyers in Table EC.4 with picking order 1, 2. The resulting allocation is boxed. It is not PO since exchanging  $a, b$  increases buyer 2’s utility while keeping 1’s unchanged. It is also not stable since  $(2, a)$  is a blocking pair.

	a	b	c	d
Buyer 1	<span style="border: 1px solid black; padding: 2px;">5</span>	5	<span style="border: 1px solid black; padding: 2px;">1</span>	1
Buyer 2	50	<span style="border: 1px solid black; padding: 2px;">1</span>	1	<span style="border: 1px solid black; padding: 2px;">1</span>