Stability, Fairness and the Pursuit of Happiness in Recommender Systems

Gerdus Benadè and Nachiketa Sahoo

Boston University

Abstract

Top-k recommendations are ubiquitous, but are they stable? We study whether, given complete information, buyers and sellers prefer to be in a platform using top-k recommendations rather than pursuing off-platform transactions among themselves. When there are no constraints on the number of exposures, we show that top-k recommendations are stable. However, stable k-recommendations may not exist when exposures are constrained, e.g., due to limited inventory or exposure opportunities. We show that maximizing total buyer welfare under unit exposure constraints is stable, Pareto optimal and swap-envy free in three restricted preference domains: orthogonal buyers, identical buyers, and buyers with dichotomous valuations. We generalize these results to arbitrary exposure constraints, formulate a polynomially-sized integer program to find stable recommendations (when they exist) and propose three other variants of common recommendation strategies adapted to satisfy exposure constraints. Experiments on three real-world datasets find these recommendation strategies exhibit substantial instability and envy. Among them, maximizing total buyers' welfare leads to the most stable outcomes.

1 Introduction

Recommender systems play an important role in market making by matching buyers to products (and 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 typically chooses one. Traditionally, recommender systems focused on satisfying individual 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 recommender systems can concentrate sales on popular sellers and increase inequities (Fleder and Hosanagar, 2009). This poses a risk to online marketplaces, 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 not only need to exhibit multi-sided fairness, but also respect different stakeholders' constraints. For example, recommending a physical good to more buyers than the number of available copies can result in a costly stockout. For platforms like ad-networks, exposure is a limited resource. This, along with contracted guarantees on exposures, limits how many times a seller can be recommended to potential buyers. Sellers, in turn, would like scarce exposures to target buyers who give them the best chance of making a sale.

Would buyers and sellers, given complete information on preferences and constraints, like to participate in a platform that uses a recommender system? Or, might some prefer to pursue off-platform transactions among themselves?

Establishing whether top-k recommendations incentivize off-platform transactions is important for two reasons. First, even in offline settings with high cost of discovery where a lack of complete information can let the market operate for some time, it has been observed[1] that unstable markets tend to unravel over time due to persistent incentives (Roth, 1984, 2015). Second, competition among digital platforms increasingly facilitates multi-homing and other forms of off-platform transactions **??**.

1.1 Our Contributions

We start with McFadden (1973)'s choice model for the behaviour of a buyer who is recommended a set of items (or choice set). We investigate three properties towards robust recommendations. *Stability* requires that both buyers and sellers would prefer to continue participating via the recommendations made by the system rather than make side-deals (Roth and Sotomayor, 1992). *Pareto optimality* (of buyer utilities) ensures efficiency, specifically, that increasing the utility of any buyer comes at the cost of another. *Envy-freeness* guarantees that no buyer prefers to receive the recommendations made to another over her own. We propose a relaxation of envy-freeness 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.

We study the existence of stable recommendations in Section 3. In the absence of exposure constraints, we observe that top-k recommendations are stable, envy free, and Pareto optimal. In the presence of exposure constraints, however, Theorem 3 shows that stable recommendations may not exist. Next, we identify three restricted preference domains in which we show that maximizing the total expected buyer utility leads to stable, Pareto optimal and SEF1 recommendations under unit exposure constraints. Notably, maximizing expected buyer utility under our choice model is equivalent to maximizing the *Nash welfare*, with respect to the exponentiated utilities, which is known to have several attractive fairness properties (Caragiannis et al., 2019). When buyers' top-k item sets satisfy the exposure constraints, recommending exactly those sets is clearly stable while maximizing buyer utility. We show in Section 3.1 that when buyers' preferences are identical, as might be the case when the quality of an item is publicly known and agreed upon, maximizing expected buyer welfare is again stable. In Section 3.2 we obtain similar positive results for buyers with dichotomous values for items, where every item is either good or bad according to buyers (and buyers may disagree on this assessment). These results generalize to arbitrary exposure constraints, with some caveats. Specifically, stability can not be guaranteed for identical buyers in general (see Example 9), though maximizing the expected buyer welfare remains PO and SEF1 (Theorem 10). In the dichotomous value setting, swap-envy freenes is impossible to guarantee, as we show in Example 12, but maximizing buyer welfare still guarantees stability and Pareto optimality (Theorem 13).

Absent the guarantee of stable recommendations, we propose three recommendation strategies in Section 4. The first algorithm draws on the theoretical results above and maximize expected buyer welfare (Section 4.1), though in general this is not guaranteed to find a stable recommendation even when one exists. Round robin recommendations are considered in Section 4.2 and shown to be SEF1 under unit exposure constraints (Theorem 15). Finally, we propose a greedy form of top-k recommendations in Section 4.3 adapted to respect general exposure constraints. This mimics decision-making in an online variant of the problem but is not guaranteed to return a recommendation that is stable, PO or SEF1.

We conclude with a computational study using datasets collected from two e-commerce platforms (Amazon and Rent-the-runway). The remaining strategies, though mirror what might be used in practice, are unstable and result in a significant incentive to deviate from the proposed matches. Among them, maximizing buyers' total utility leads to the best outcomes.

[1] When matching interns to hospitals and matching children to schools.

Our theoretical and computational results suggest that though stable recommendation may not always be possible, maximizing buyers' welfare often achieves a good trade-off between keeping buyers and sellers happy.

1.2 Related Work

In the recommender system literature, multi-stakeholder recommendations which consider all buyer, seller and platform preferences are increasingly popular (Burke et al., 2016; Nguyen et al., 2017; Abdollahpouri et al., 2020). Fleder and Hosanagar (2009) find that buyer-focused recommender systems can concentrate sales on popular sellers and increase inequities among sellers. In contrast, we identify several settings where maximizing buyer welfare is aligned with the stability of the market. The study of fairness in recommender systems is relatively recent and is sometimes in the context of multi-objective recommendations (Ekstrand et al., 2022; Abdollahpouri and Burke, 2019; Ziegler et al., 2005; Patro et al., 2020). Bateni et al. (2022) propose a stochastic approximation scheme based on the Eisenberg-Gale convex program for online advertising system which maximizes platform revenue while being approximately fair towards buyers. We similarly draw inspiration from the Eisenberg-Gale program to find a setting where the interests of different stakeholders (buyers and sellers, in our case) are aligned. Like us, Patro et al. (2020) view fair recommendation as a fair allocation problem. They propose greedy version of round robin which is shown to be envy-free up to one good (EF1) for buyers and guarantees sellers some minimum level of exposure. We consider an item's purchase probability as metric subject to satisfying exposure constraints rather than a minimum exposure level and go beyond envy by additionally seeking stable recommendations.

There is a long literature on stable matchings dating back to the 1950's (Stalnaker, 1953; Gale and Shapley, 1962), for more thorough treatments we refer the interested reader to Roth and Sotomayor (1990) or Abdulkadiroglu and Sönmez (2013). Recommending k items to a buyer reminds strongly of worker-firm (Kelso and Crawford, 1982) or college admissions (Gale and Shapley, 1962) matching programs where workers (students) are matched to firms (colleges), sometimes subject to quotas on the number of matches. Kelso and Crawford (1982) show that when preferences satisfy a substitutability condition and when workers' preferences depend only on the firm they apply to, not their co-workers, then stable many to one matchings exist. Our setting does not satisfy these conditions, as the purchase probability of an item depends directly on the other items recommended to the buyer. Indeed, items (and their sellers) have preference orders not over buyers but over (buyer, choice set) pairs. This leads to complications even with unit exposure constraints. In traditional deferred acceptance schemes, one side of the market proposes matches to the other in order of their preferences, and matches are tentatively accepted until a better proposal comes along. In our setting, the purchase probability of an item depends on all other items recommended to the same seller. Hence, the seller of an item can not accurately judge the attractiveness of a buyer's offer until the other $k-1$ items in that buyer's choice set are fixed.

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., 2016) in both static and dynamic settings (Benadè et al., 2018; Zeng and Psomas, 2020). The concept of Nash welfare, or the product of agent utilities, originated in John Nash's solution to a bargaining problem (Nash, 1950). Maxmimizing Nash welfare when dividing indivisible goods among agents with additive utilities is known to be EF1 and Pareto optimal (PO) (Caragiannis et al., 2016). Maximizing Nash welfare is NP-hard for several bidding languages (Ramezani and Endriss, 2010); Caragiannis et al. (2016) propose a computational approach which scales to reasonably sized instances. In a similar spirit to our work, 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 4.2 and further show that maximizing expected buyer welfare is also SEF1. Igarashi et al. (2022) study the fair allocation of players to teams where teams have additive values for players and players have a weak preference order over teams. It is found that maximizing Nash welfare is EF1 and PO, as it is in the traditional one-sided case. Igarashi et al. (2022) also consider two stability notions. An allocation is said to be *swap stable* when there is no pair of teams and players on those 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. Our notion of stability requires only that the deviating buyer and seller are strictly better off, not that the other parties involved are no worse off.

2 Model

We study a setting in which each of a batch of buyers is simultaneously recommended k items, each item subject to constraints on the number of times it is recommended. Let $\mathcal{B}$ denote a set of n buyers and $\mathcal{I}$ a set of m items. For simplicity we assume every item is sold by a different seller, so we may occasionally refer to recommending sellers, rather than items, to buyers.

A *(k−)recommendation* to buyer $b \in \mathcal{B}$ is a set of k items $\bar{A}_b \subseteq \mathcal{I}$, $|\bar{A}_b| = k$. In addition to $\bar{A}_b$, buyer b has (fixed) outside option ω_b available which represents not selecting any of the recommended items and instead sticking with the status quo or pursuing an off-platform transaction. Let $A_b = \bar{A}_b \cup \{\omega_b\}$. We call the vector of recommendations $A = (A_b)_{b\in\mathcal{B}}$ a *recommendation profile*. Let $A_i^{-1} = \{b \in \mathcal{B} : i \in A_b\}$ denote the buyers recommended item i. A recommendation profile is feasible if it satisfies constraints on the number of exposures received by each item, encoded as $|A_i^{-1}| = c_i$ for all $i \in \mathcal{I}$. Let $\mathcal{A}$ denote the set of feasible recommendation profiles.

We occasionally restrict our analysis to the unconstrained ($c_i = \infty$ for all $i \in \mathcal{I}$) and unit constrained ($c_i = 1$ for all $i \in \mathcal{I}$) settings. Under unit exposure constraints each item is recommended to only one buyer. We overload notation and let A_i^{-1} refer to the buyer recommended item $i \in \mathcal{I}$ and generally use i interchangeably with the singleton set $\{i\}$.

Buyer behavior is assumed to follow a standard choice model. Buyer $b \in \mathcal{B}$ has utility $V_{bi} = v_{bi} + \epsilon_{bi}$ for item $i \in \mathcal{I}$, where v_{bi} is the observable *value* that b has for i and ϵ_{bi}, drawn independently and identically from a Gumbel distribution,[2] is a random utility component which is unknown in advance and captures unobserved determinants of item utility. We take values v_{bi} as arbitrary and known, for example, it may have been estimated from incomplete data using a collaborative filter. Buyer b has value $v_{b\omega}$ for outside option ω_b, which reflects the utility derived from pursuing an off-platform transaction. For ease of exposition we assume all buyers have $v_{b\omega} = v_\omega$. This assumption is not critical: when buyers have outside options of different quality we can denote with $v'_{bi} = v_{bi} - v_{b\omega}$ the normalized utility of item i — all our results remain true in terms of normalized utilities.

Per choice model theory, buyer $b \in B$ considering choice set S realises the unknown random components of their utility functions, then deterministically selects the option $i \in S$ that provides greatest utility. At the time of recommendation, a buyer's *welfare*, her expected utility from the entire choice set, is given by $v_b(S) = \mathbb{E}\left(\max_{i\in S}(V_{bi})\right) = \log\left(\sum_{i\in S} e^{v_{bi}}\right)$ (Williams, 1977). The *value profile* associated with recommendation A is $v(A) = (v_b(A_b))_{b\in\mathcal{B}}$. For convenience we call $u_{bi} = e^{v_{bi}}$ the *virtual value* of buyer b for item i, set $u_b(S) = \sum_{i\in S} e^{v_{bi}}$ and call $u(A) = (u_b(A_b))_{b\in\mathcal{B}}$ the *virtual value profile*. Notice that $v_b(S) = \log u_b(S)$.

It has been shown that, at the time of recommendation, the probability that option i will be selected from choice set S is $\mathbb{P}(b, i, S) = e^{v_{bi}} / \sum_{j\in S} e^{v_{bj}}$. Under unit exposures, the probability of item $i \in \mathcal{I}$ being purchased under A is $\mathbb{P}(i, A) = \mathbb{P}(A_i^{-1}, i, A_{A_i^{-1}})$, and generally $\mathbb{P}(s, A) = \sum_{b\in A_i^{-1}} \mathbb{P}(b, i, A_b)$. The welfare of (the seller of) item i is

[2] Also called a type I extreme value distribution.

assumed only to be increasing in $\mathbb{P}(i, A)$, which is flexible enough to accommodate different sellers having different profit margins.

Notice that buyers have a fixed preference order over $\mathcal{I} \cup \{\omega\}$, as determined by $\{v_{bi} : \forall i \in \mathcal{I}\{\omega\}\}$, regardless their choice set. However, sellers of items do not have a fixed preference order over buyers: the probability of a sale depends not only on how much the buyer values the item, but also on the total virtual value of other items in that buyer's choice set.

2.1 Measuring The Quality of a Recommendation

Efficiency One natural requirement is that the recommendation profile is *Pareto optimal* (PO) with respect to buyers' welfare. Let $[n] = \{1, \ldots, n\}$. Formally, 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 $A \in \mathcal{A}$ is Pareto optimal if there does not exist another recommendation $A' \in \mathcal{A}$ such that $v(A')$ strictly dominates $v(A)$. Notice that $v(A)$ is undominated exactly when $u(A)$ is undominated.

Fairness A standard notion of fairness is called *envy-freeness*, which we study from the buyers' perspective. Envy-freeness requires that every buyer prefers their choice set over the choice set of any other buyer. Formally, recommendation profile A is envy-free when $v_b(A_b) \geq v_b(A_{b'})$ (equivalently, $u_b(A_b) \geq u_b(A_{b'})$) for all $b, b' \in \mathcal{B}$. Envy-freeness is often impossible with indivisible objects — consider allocating a single valuable item to two agents. As a result, envy-freeness is commonly relaxed to envy-freeness up to one item (EF1) for indivisible goods, which allows envy 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 new relaxation of envy-freeness, called *swap-envy-freeness*, to capture this. We say 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 there exist a pair of items $i \in A_b, j' \in A_{b'}$ so that $v_b(A_b \cup j \setminus i) \geq v_b(A_{b'} \cup i \setminus j)$.

Stability Pareto optimality and (swap-)envy-freeness both consider only buyer welfare. However, they do not guarantee that buyers and sellers won't have incentives to deviate from a recommendation profile. 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 from b replacing some item j in A_b with i. Formally, (b, i) blocks when $i \notin A_b$ and there exists a seller $j \in A_b$ so that both $v_{bi} > v_{bj}$ and $\mathbb{P}(i, A) < \mathbb{P}(b, i, A_b \cup i \setminus j)$. To handle general exposure constraints and the possibility that $j \in A_c$, we further require that the recommendation profile resulting from the deviation can be feasibly extended to a full recommendation. This condition has no bite for the unconstrained and unit exposure settings, and is elsewhere equivalent to assuming that there are dummy items available for which ever buyer has the lowest possible value. A recommendation profile is *stable* in the absence of a blocking pair.

We illustrate some differences between these concepts with an example.

Example 1. *Consider an instance with buyers {1,2}, four items $\{a, b, c, d\}$ and $k = 2$. Table 1 shows the virtual values for an instance with identical buyers and two good and two bad items.*

Let's assume each item can be recommended only once. 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, although buyer 2 envies buyer 1: 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.

	a	b	c	d
Buyer 1	2	2	1	1
Buyer 2	2	2	1	1

Table 1. Buyers' virtual values.

3 On The Existence of Stable Recommendations

We investigate the conditions under which stable recommendations exist. When the exposure constraints are not binding, there is no need to do anything more complicated than recommending each buyer their k highest value items.

Observation 2. *In any instance where the exposure constraints are not binding, for example when $c_i = \infty$ for all $i \in \mathcal{I}$, or 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, is a direct result of the fact that a buyer who is recommended their k highest value items will be unwilling to participate in a blocking pair. To the best of our knowledge, this property of top-k recommendations has not been discussed before. It is worth pointing out because capacity constraints generally do not exist for digital products, a category of goods in which recommendation systems are often deployed, so it is reassuring to know that top-k recommendation is stable.

Unfortunately, binding exposure constraints preclude the existence of stable recommendations.

Theorem 3. *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.

	a	b	c	d
Buyer 1	10	0	7	6
Buyer 2	10	8	4	5

Table 2. Buyers' virtual values for each seller.

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 0 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. □

The fact that stability can not be guaranteed should not come entirely as a surprise. For example, Pycia (2012) show that stable coalitions can only be guaranteed when over any two coalitions all common members prefer the same coalition. In our setting, recommendations are the coalitions of a buyer and k sellers. If the buyer obtains higher welfare by changing some items of the choice set, the purchase probability of the remaining items must reduce due to increased competition. As a result, buyers and sellers of remaining items prefer opposing choice sets.

A closer look at the instance in Theorem 3 reveals two situations which may lead to the existence of blocking pairs. First, whenever b was allocated to buyer 1, the massive discrepancy in in the buyers' values for item b created the potential for b to increase its purchase probability by deviating from the recommendations. Second, the allocations with $b \in A_2$ lead to a large disparity in the bundle values, specifically, $|u_1(A_1) - u_2(A_2)| \geq 4$. When the choice sets are not equally competitive even an item valued identically by the buyers may participate in a blocking pair, 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 degree to which buyers value the same item similarly and the difference in competition across choice sets.

Definition 4 (α-balance)**.** *The valuations of an item i is called α_i-balanced if $u_{bi} \leq \alpha_i \cdot u_{ci}$ for all $b, c \in \mathcal{B}$. When each $i \in \mathcal{I}$ is α_i-balanced, the instance is called α-balanced, with $\alpha = \max_i \alpha_i$.*

When buyers have identical values the instance is 1-balanced. Next we parameterize the level of competition across choice sets.

Definition 5 (β-impartiality)**.** *An allocation A is called β-impartial when β is the smallest value such that $u_b(A_b) \leq \beta \cdot u_c(A_c)$ for all $b, c \in \mathcal{B}$.* [3]

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 $A'_c = A_c \cup j \setminus i$, in other words, 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 6. *Consider an α-balanced instance with β-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 upper bound as*

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

As an example, consider allocation $A_1 = \{a, c\}, A_2 = \{b, d\}$ in the instance of Theorem 3 which has $\alpha_a = 1$ and $\beta = \frac{17}{13}$, implying a maximum benefit of $\frac{17}{13} \approx 1.3$. Participating in the blocking pair $(2, a)$ increases seller a's purchase probability by a factor of $\frac{10/15}{10/17} \approx 1.13$, while buyer 2's welfare increases by a factor of $\frac{15}{13} \approx 1.15$.

In the remainder of this section we study two other restricted preference domains, identical buyers and dichotomous values, where it is possible to guarantee stability. Omitted proofs may be found in the appendix.

3.1 Identical Buyers

Suppose buyers agree on a common evaluation v_i of each item i, i.e. $v_i = v_{bi}$ for all $b \in \mathcal{B}, i \in \mathcal{I}$.

Let A^* be a recommendation profile which maximizes total buyer welfare.[4] In other words, A^* is a solution of

$$\max_{A \in \mathcal{A}} \left\{ \sum_{b \in \mathcal{B}} \log \left(\sum_{i \in A_b} e^{v_{bi}} \right) \right\} \equiv \max_{A \in \mathcal{A}} \left\{ \sum_{b \in \mathcal{B}} \log(u_b(A_b)) \right\} \equiv \max_{A \in \mathcal{A}} \left\{ \prod_{b \in \mathcal{B}} u(A_b) \right\}.$$

Notice that maximizing buyer welfare equivalently maximizes the *Nash welfare* (product of utilities) with respect to the virtual values. In the fair division literature, maximizing Nash welfare (on item values) is known to be EF1 and PO (Caragiannis et al., 2019). It also works here in the case of unit exposure constraints.

[3] In a similar spirit to α_i-balancedness above one could consider consider β_b-impartiality, where β_b is the smallest value so that $u_b(A_b) \leq \beta_b u_c(A_c)$ for all $c \in \mathcal{B}$.

[4] Incidentally, recommending each buyer her top-k maximizes the total buyer welfare when exposure constraints are not restrictive.

Theorem 7. *For buyers with identical preferences under unit exposure constraints, A^* is stable, PO and SEF1.*

We rely on the following technical lemma, which is straightforward to verify.

Lemma 8. *For all $x, y \geq 0$ satisfying $x + y = C$ for a fixed constant C, xy is monotonically increasing in* $\min\{x, y\}$.

Proof sketch for Theorem 7. By Lemma 8, maximizing the product of agent virtual values will encourage the agent virtual values to be as similar as possible. Since buyers have identical values, a seller s currently allocated to buyer b will only participate in a blocking pair (c, s) if she faces less competition post transfer. This would imply buyer c's value for their allocation (post-deviation) is lower than b's value for their original allocation. We show this contradicts A^* being an optimal solution by lemma 8, stability follows. SEF1 is a result of agents having similar bundle virtual values and the monotonicity of the logarithm: if the allocation was not SEF1 there would exist a swap which increases the objective function (by lemma 8). Pareto optimality follows from A^* being a maximizer. □

Under unit exposure constraints, maximizing the product of virtual values leads to choice sets that are competitive enough to prevent profitable deviations. As the following example shows, general exposure limits can lead to less competitive choice set. As a result, maximizing buyer welfare is not guaranteed to be stable.

Example 9. *Consider the following instance with $n = 2, k = 3$ and six unique items and values and capacities as in Table 3. We may assume buyer 1 receives the most valuable item (the other case is symmetric). The resulting recommendation which maximizes the product of buyers' virtual values (equivalently, expected welfare) is indicated with checkmarks.*

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 19 to 21. Item a's purchase probability also increased, from $\frac{12}{25} < 0.5$ to $\frac{12}{21} > 0.5$. We conclude that $(2, a)$ is a blocking pair. As such, the recommendation found by maximizing buyer welfare is not stable.

Notice that in the corresponding instance with seven items (two distinct items with value 10) and unit constraints the product maximizing choice sets are more competitive with virtual values 21 and 23, rather than 19 and 25.

	a	b	c	d	e	f
Values	12	10	5	4	3	3
Capacity	1	2	1	1	1	1
Buyer 1	✓	$\boxed{✓}$			$\boxed{✓}$	$\boxed{}$
Buyer 2	$\boxed{}$	✓	$\boxed{✓}$	$\boxed{✓}$		

Table 3. Buyers' virtual values.

Despite stability breaking down, A^* remains PO and SEF1. The argument for SEF1 requires showing that whenever a buyer envies another, there exists a feasible exchange of items which decreases that envy. Because buyers have identical values, it is possible to show that any swap decreasing envy either eliminates the envy completely or increases the product of virtual values.

Theorem 10. *Under general exposure constraints, A^* is PO and SEF1 for buyers with identical preferences.*

3.2 Dichotomous Values

Consider the case where buyers have dichotomous preferences and has either high or low value for every item (and is indifferent among items of each type). Formally, assume $v_{bi} \in \{a, a'\}$ for all $b \in \mathcal{B}$ and $i \in \mathcal{I}$, for some real-valued

$a < a'$. Note that buyers need not agree on the evaluation of an item: it may be acceptable to some and unacceptable to others. Nor do they need to prefer disjoint sets of items. In this setting, maximizing the total buyer welfare (or maximizing Nash welfare with respect to the virtual values) again leads to stability.

Theorem 11. *For buyers with dichotomous values, A^* is stable, PO and SEF1 under unit exposure constraints.*

Proof. To ease exposition, 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, PO and SEF1.

Stability: Assume, for contradition, 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 participate in a blocking pair. Let $j \in A^*_b$ be such an item with $u_{bj} = \ell$. Construct A' from A^* by exchanging i and j. So $A'_b = A^*_b + i - j$, $A'_c = A^*_c + j - i$ and $A'_d = A^*_d$ for all $d \in \mathcal{B} \setminus \{b, c\}$.

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

1. (ℓ, ℓ): Participating in a blocking pair requires displacing a buyer with strictly lower value, which is not possible since $u_{bi} = \ell$. This contradicts (b, i) being a blocking pair.

2. (ℓ, h): Again, $u_{bi} = \ell$, which contradicts (b, i) being a blocking pair.

3. (h, ℓ): In this case i is assigned to c despite c having low and b having high value for it. It follows that $u_b(A') = u_b(A^*) - \ell + h > u_b(A^*)$ and $u_c(A') = u_c(A^*) - \ell + u_c(j) \geq u_c(A^*)$. However, this implies $u_b(A') \cdot u_c(A') > u_b(A^*) \cdot u_c(A^*)$, contradicting the fact that A^* maximizes the product of virtual values.

4. (h, h): Suppose that $u_{cj} = h$. Now $u_b(A') > u_b(A^*)$ and $u_c(A') = u_c(A^*)$, contradicting that A^* maximizes the product of virtual values. We conclude that $u_{cj} = \ell$ and thus $u_b(A') + u_c(A') = u_b(A^*) + u_c(A^*)$, so we can apply Lemma 8.

   We now compare $\min\{u_b(A^*_b), u_c(A^*_c)\}$ and $\min\{u_b(A'_b), u_c(A'_c)\}$. A condition for (b, i) being a blocking pair is that i increases their purchase probability by deviating, which implies $u_{bi}/u_b(A'_b) > u_{ci}/u_c(A^*_c)$. Since $u_{bi} = u_{ci} = h$, it follows that $u_b(A'_b) < u_c(A^*_c)$. Similarly, b improves their utility by participating in the blocking pair, so $u_b(A^*_b) < u_b(A'_b)$. It follows that $u_b(A^*_b) = \min\{u_b(A^*_b), u_c(A^*_c)\}$. It remains to compare $u_b(A^*_b)$ and $u_c(A'_c)$. Now

   $$u_c(A'_c) = u_c(A^*_c) + \ell - h > u_b(A'_b) + \ell - h = (u_b(A^*_b) + h - \ell) + \ell - h = u_b(A^*_b)$$

   which implies $u_b(A^*_b) < \min\{u_b(A'_b), u_c(A'_c)\}$. By lemma 8, this contradicts A^* maximizing the product of virtual values.

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.

$$u_b(A^*_b) + u_{bi} - u_{bj} < u_b(A^*_c) - u_{bi} + u_{bj} \text{ for all } j \in A^*_b, i \in A^*_c.$$

Consider arbitrary $j \in A^*_b$ and $i \in A^*_c$ so that $u_{bj} = \ell$ and $u_{bi} = h$ (such items exist since b envies c). If it were the case that $u_{ci} = \ell$, 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 maximimality of A^* thus implies that $u_{ci} = h$. We similarly conclude that $u_{cj} = \ell$. This means both buyers agree that every item $j \in A_b^*$ with $u_{bj} = \ell$ is a low value item, and similarly every $i \in A_c^*$ with $u_{bj} = h$ is a high value item.

Let $H_b = \{j : u_{bj} = h\}$ and $L_b = \{j : u_{bj} = \ell\}$, 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 $A'_b = A_b^* + i - j, A'_c = A_c^* + j - i$ and $A'_d = A_d^*$ for all $d \in \mathcal{B} \setminus \{b, c\}$. Notice that, since b and c values i and j identically, $u_b(A_b^*) = u_c(A_c^*) = u_b(A'_b) + u_c(A'_c)$ so we are in a position to apply Lemma 8 if appropriate.

Since A^* is not SEF1, $u_b(A'_b) < u_b(A'_c)$. Buyer c receives an ℓ item in exchange for an h one, so $u_c(A'_c) < u_c(A^*)$, similarly $u_b(alloc'_b) > u_b(A_b^*)$. Because c also has high value for those items in A'_c that b has high value for,

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

Putting it all together shows $u_b(A_b^*) < u_b(A'_b) < u_b(A'_c) \leq u_c(A'_c) < u_c(A_c^*)$, so exchanging i and j increases the minimum utility. By Lemma 8, this contradicts A^* maximizing the product of virtual values. We conclude that A^* is SEF1 in terms of virtual values. SEF1 in terms 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 one buyer values the commonly recommended items much lower than other and the buyers agree on which of the other items are good or bad. Maximizing buyer welfare will result in allocating the good items to the buyer who values the common set least, and the bad the 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 12. *Consider the following instance with $n = 2, k = 15$ and 20 unique items of three different 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 1, which both buyers value at 1. The instance is summarised in Table 4*

The unique recommendation which maximizes the product of buyers' virtual values (equivalently, expected welfare) is boxed in Table 4: 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 25 for her own choice set and 30 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 present when converting the virtual values into expected utilities.

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 13. *Under general exposure constraints, A^* is stable and PO for buyers with dichotomous values.*

Item type	a	b	c
# of items	10	5	5
Capacity	2	1	1
Buyer 1	$\boxed{2}$	2	$\boxed{1}$
Buyer 2	$\boxed{1}$	$\boxed{2}$	1

Table 4. Buyers' virtual values.

3.3 Other Restricted Preference Domains

We briefly remark on two other common structured preference domains. First, suppose buyers have factored utility structure: each buyer b is associated with a characteristic vector β_b, each item i with similar vector γ_i, and $v_{bi} = \langle \beta_b, \gamma_i \rangle$. A stable recommendation need not exist in setting: the instance in the proof of Theorem 3 can be (approximately) 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 existence is not guaranteed for larger instances. Full details appear in the appendix.

4 Recommendation Strategies Under Exposure Constraints

Despite it not being possible to guarantee stability for general instances, one my still attempt to find a stable recommendation on those instances that permit it. In Appendix A.4 we construct a polynomially-sized integer program which find stable recommendations when they exist and otherwise minimizes the largest incentive any seller has for participating in a blocking pair. Unfortunately, this approach does not appear to scale to realistic instances.

In this section we propose three alternative practical recommendation strategies that can accommodate constraints on number of exposures per product. Throughout the remainder of this paper we focus on the case of unit constraints, but all strategies are simple to modify to arbitrary capacity constraints.

4.1 Maximizing Total Buyer Welfare

A common platform objective is to maximize the total buyer welfare, which was shown to have the additional benefit of leading to stable, fair and efficient recommendations in specific preference domains. The resulting problem is

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

where $x_{bi} = 1$ when i appears in b's choice set. This is equivalent to maximizing the Nash welfare (product of utilities) with respect to virtual values $u_{bi} = e^{\hat{v}_{bi}}$, which is NP-hard (Nguyen et al., 2013).

Because the total buyer welfare is directly maximized, buyers may be less willing to cooperate in deviations. However, even for unit capacities maximizing buyer welfare is not guaranteed to find a stable recommendation in instances where one exists.

Example 14. *Consider the instance in Table 5. The recommendation that maximizes buyer welfare is $A_1 = \{a, d\}, A_2 = \{b, c\}$ (boxed in Table 5), however, $(1, c)$ is a blocking pair. At the same time, the underlined recommendation $A_1 = \{c, d\}, A_2 = \{a, d\}$ is stable.*

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

Table 5. Buyers' virtual values.

4.2 Round Robin Recommendations

This strategy fixes a permutation of buyers, then cycles through the buyers k times. At each step, the active buyer is assigned their highest value item from the remaining items.

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 over greedy top-k, is that common high value items are shared among buyers. This may increase the probability of purchase for high value items by avoiding hyper-competitive choice sets. That round robin leads to approximately equal welfare across buyers further suggests that all choice sets will be similarly competitive. This reinforces the expectation that sellers may have less incentive to deviate, which would make stability more likely.

We can show that round robin, which is EF1 in general settings, is SEF1, though it need not be stable or PO.

Theorem 15. *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 every exchange of items between their allocations. We may assume without loss of generality that b was after c in the permutation of buyers, otherwise b would not envy c. We ignore the other buyers, since their bundles do not affect b's envy towards c.

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

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

Let $j = \arg\max_{j\in[k]} u_b(c_j)$ be the index of the item in A_c that b likes most. Create a new allocation A' by swapping b_k and c_j and leaving the other buyers unchanged, so $A'_b = A_b - b_k + c_j$, $A'_c = A_c + b_k - c_j$ and $A'_d = A_d$ for all $d \in \mathcal{B} \setminus \{b, c\}$.

We will proceed to match items in A'_b and A'_c so that b always likes the item they received more that the corresponding item that c received. The existence of such a matching clearly implies that b does not envy c after summing over all pairs. Consider the following cases:

- $j = 1$: $A'_b = \{c_1, b_1, \ldots, b_{k-1}\}$ and $A'_c = \{b_k, c_2, \ldots, c_k\}$. By assumption, $u_b(c_1) > u_b(b_k)$, so we may match c_1 and b_k. We can also match b_i and c_{i+1} for all $i \in [k-1]$, since we know $u_b(b_i) > u_b(c_{i+1})$ by the nature of round robin. We conclude that b does not envy c in A'.
- $j > 1$: $A'_b = \{c_j, b_1, \ldots, b_{k-1}\}$ and $A'_c = \{c_1, \ldots, c_{j-1}, c_{j+1}, \ldots, c_k, b_k\}$ (possibly with $1 = j-1$ or $j+1 = k$). By choice of j, $u_b(c_j) > u_b(c_1)$, so match c_j and c_1. For all $i \in [k-1] \setminus \{j-1\}$, match b_i and

c_{i+1} as before. Finally, note that $u_b(b_{j-1}) > u_b(b_k)$ since it was assigned to b earlier during the round robin procedure, so we can match b_{j-1} and b_k and conclude that b does not envy c in A'.

In either case, we arrive at a contradiction. We conclude that round robin is SEF1. □

4.3 Greedy Top-k

This strategy iterates over buyers in a random order and assigns each buyer the choice set of size k which would maximize her welfare. This mimics how recommendations are typically made online. Buyers arrive on a website one at a time, and the site recommends the set of items that would provide the buyer the highest utility without considering future arrivals. 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 $k \geq 4$ 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 $[k]$ at 1 and the remainder at 0. Because buyer 1 is considered first, without regard for other buyers, buyer 2 is recommended items $\{k+1, \ldots, 2k\}$ which they have 0 value for. The resulting recommendation is clearly neither Pareto optimal nor SEF1.

Proposition 16. *Greedy top-k recommendations need not be stable, SEF1 or Pareto optimal.*

5 Measuring Instability in Real World Datasets

Theorem 3 shows that there are instances where stability is impossible, but it may be the case that such instances are extremely rare. Moreover, even in the absence of stability it is possible that there are very few sellers who 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 recommendations in real-world datasets and, if not, whether sellers have a significant incentive to deviate from the recommendation. We do this by predicting ratings ($\hat{v}_{bi}$) for all buyer–item pairs and deploying recommender systems in three real world datasets. Throughout this section we assume unit exposure constraints.

5.1 Datasets

We use three datasets for our experiments: two containing customers' ratings on products from Amazon (in the Automotive and Musical Instruments categories) and one containing 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 directly 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.

				Rating summary			
Dataset	# users	# items	# ratings	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.

5.2 Experiments

We simulate the following scenario. On any given day a subset of the buyers visit the platform. The platform recommends $k = 5$ items to each buyer. For simplicity, we assume that the platform has only one copy of each item in its inventory and does not recommend any item to more than one buyer to avoid the risk of stockout.

We train SVD++, a matrix factorization based collaborative filter, on the entire dataset (Koren, 2008).[5] 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 value for the item. We track five metrics for each recommendation strategy:

1. **Move**: the percentage of sellers who participate in a blocking pair, i.e. are able and prefer to move to another buyer. This measures how *widespread* the incentive to deviate is.

2. **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.

3. **Welfare**: the average buyer welfare from their recommended choice sets. All else being equal, a platform would like to provide its buyers higher welfare.

4. **Envy**: the percentage of buyers who envy another. Lower envy may signal that the recommendation treats buyers more equally.

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 (it can be removed by a single swap).

We draw 16 random buyer-item pools for each dataset and report average metrics in Table 7.

There is widespread and substantial incentive to deviate from the matchings formed by top-k recommendation under each of these strategies. Greedy top-k is particularly problematic, 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 match.[6] 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; under this strategy roughly half the sellers have incentive to deviate but they stand to gain only about 10%. Round robin performs a bit worse, but the simplicity of the algorithm may make it an attractive option in the event that it is not computationally feasible to maximize utility on large datasets.

[5] 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.

[6] 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 under this setting, since particularly worse off buyers will attract multiple new sellers, thereby leading to increased competition.

Table 7. Instability in top-5 recommendations under unit exposure constraint.

	Number of users/items									
	50/250					200/1000				
Automotive	Move $^-$	Gain $^-$	Welfare $^+$	Envy $^-$	Swap Envy $^-$	Move $^-$	Gain $^-$	Welfare $^+$	Envy $^-$	Swap Envy $^-$
Greedy top-5	97.08%	114.10%	6.27	76.25%	59.62%	99.12%	141.50%	6.36	71.28%	58.53%
	(0.22%)	(4.50%)	(0.008)	(1.82%)	(1.14%)	(0.07%)	(3.03%)	(0.006)	(0.86%)	(0.92%)
Round robin	64.75%	11.50%	6.36	41.75%	0.00%	71.06%	11.69%	6.44	38.31%	0.00%
	(1.06%)	(0.30%)	(0.007)	(1.00%)	(0.00%)	(0.46%)	(0.16%)	(0.006)	(0.72%)	(0.00%)
Max. welfare	50.95%	8.59%	6.39	47.38%	0.13%	57.19%	9.06%	6.47	46.91%	0.13%
	(1.04%)	(0.36%)	(0.007)	(2.61%)	(0.12%)	(0.73%)	(0.25%)	(0.005)	(1.43%)	(0.08%)
Musical Instr.										
Greedy top-5	97.08%	99.12%	6.26	74.75%	58.25%	99.16%	131.50%	6.33	73.03%	59.16%
	(0.26%)	(5.55%)	(0.008)	(1.45%)	(1.65%)	(0.06%)	(3.82%)	(0.005)	(0.75%)	(0.98%)
Round robin	64.65%	10.50%	6.34	39.62%	0.00%	69.42%	10.70%	6.40	36.44%	0.00%
	(0.88%)	(0.23%)	(0.007)	(1.58%)	(0.00%)	(0.33%)	(0.17%)	(0.005)	(0.55%)	(0.00%)
Max. total utility	49.72%	8.01%	6.37	46.87%	0.50%	55.66%	8.63%	6.43	52.00%	1.13%
	(1.12%)	(0.28%)	(0.007)	(3.25%)	(0.28%)	(0.89%)	(0.29%)	(0.005)	(1.84%)	(0.29%)
Rent-the-runway										
Greedy top-5	97.57%	147.70%	10.88	89.00%	78.87%	99.17%	174.70%	10.95	88.44%	78.44%
	(0.10%)	(5.64%)	(0.008)	(1.30%)	(1.43%)	(0.03%)	(4.10%)	(0.005)	(0.37%)	(0.50%)
Round robin	69.20%	14.21%	11.00	52.88%	0.00%	73.56%	14.43%	11.07	50.50%	0.00%
	(0.91%)	(0.11%)	(0.009)	(2.08%)	(0.00%)	(0.38%)	(0.18%)	(0.004)	(1.18%)	(0.00%)
Max. welfare	50.78%	9.76%	11.04	52.50%	0.00%	60.68%	11.61%	11.11	51.75%	0.28%
	(1.66%)	(0.42%)	(0.009)	(2.66%)	(0.00%)	(1.38%)	(0.47%)	(0.004)	(1.88%)	(0.09%)

Note: 1) Move: The % of all the sellers who can and would like to move to a different buyer. 2) Gain: The % they would gain in probability of purchase by doing so (over a baseline average of ≈ 0.2). 3) Utility: The welfare of the buyers from their choice sets. 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 due to space constraints. $^-$: smaller the better; $^+$: larger the better.

Unsurprisingly, maximizing total welfare leads to the highest buyer utility. Round robin is not far behind and always yields utility within 1% of optimal. Greedy top-k's myopic decisions are punished by the fact that there are decreasing marginal utility to adding items to the choice set.

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 and max welfare do much better, though roughly 50% of buyers still have envy. Notably, the recommendations of round robin and max welfare are almost always SEF1, meaning whatever envy exists can be eliminated by exchanging a single pair of items between choice sets. Round robin has lowest envy and swap envy in the the Amazon datasets, despite the simplicity of the algorithm.

6 Summary

Recommender systems play an important role in matching buyers to sellers in large marketplaces. Given full information, do buyers and sellers have an incentive to continue participating in such systems? Or, might some prefer to pursue alternative matches, which may unravel the market?

We observe that when there is no constraint on the number of times an item can be recommended, top-k recommendations are stable: buyers and sellers do not have an incentive to deviate from the centralized matches. However, in the presence of constraints on the number of exposures, stable k-recommendations may not exist. We show that when buyers have identical utility functions, disjoint sets of top-k items, or dichotomous utilities for all items, stable matches do exist and can be found by selecting choice sets to maximize the total buyer welfare.

In computational experiments on data collected from three markets we find that natural recommender systems exhibit significant instability. Maximizing the total buyer welfare, which is guaranteed to lead to stable outcomes in the three restricted settings, also leads to more stable outcomes than greedily recommending each buyer her top-k

or allocating items to buyers in a round-robin fashion in general. Finally, an analysis of buyers' envy shows that greedy top-k strategy generates substantial envy that is not eliminated even after swapping an item with the envied. Maximizing total buyer welfare and round robin though lead to lower, but still substantial, share of envious buyers. However, they are almost entirely eliminated upon swapping a single item.

In sum, our results suggest that while stable recommendation need not always exist, maximizing buyers welfare often achieves a good trade-off between keeping buyers and sellers happy. Characterizing exactly when stability can be guaranteed, and checking whether real-world datasets satisfy the conditions, remain as important open problems.

References

H. Abdollahpouri and R. Burke. Multi-stakeholder recommendation and its connection to multi-sided fairness. *arXiv preprint arXiv:1907.13158*, 2019.

H. Abdollahpouri, G. Adomavicius, R. Burke, I. Guy, D. Jannach, T. Kamishima, J. Krasnodebski, and L. Pizzato. Multistakeholder recommendation: Survey and research directions. *User Modeling and User-Adapted Interaction*, 30(1):127–158, 2020. doi: 10.1007/s11257-019-09256-1.

A. Abdulkadiroglu and T. Sönmez. Matching markets: Theory and practice. *Advances in Economics and Econometrics*, 1:3–47, 2013.

A. Alkan, G. Demange, and D. Gale. Fair allocation of indivisible goods and criteria of justice. *Econometrica*, 59(4): 1023–1039, 1991.

M. Bateni, Y. Chen, D. F. Ciocan, and V. Mirrokni. Fair resource allocation in a volatile marketplace. *Operations Research*, 70(1):288–308, 2022.

J. G. Benadè, A. Kazachkov, C.-A. Psomas, and A. D. Procaccia. How to make envy vanish over time. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 593–610, 2018.

S. J. Brams and A. D. Taylor. An envy-free cake division protocol. *The American Mathematical Monthly*, 102(1): 9–18, 1995.

R. D. Burke, H. Abdollahpouri, B. Mobasher, and T. Gupta. Towards multi-stakeholder utility evaluation of recommender systems. *UMAP (Extended Proceedings)*, 750, 2016.

I. Caragiannis and S. Narang. Repeatedly matching items to agents fairly and efficiently. *arXiv preprint arXiv:2207.01589*, 2022.

I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash product. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 305–322, 2016.

I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–32, 2019.

L. Chang. Rent the Runway is launching a new way for designers to partner with the company — and it's great news for subscribers, 11 2018. URL https://www.insider.com/guides/style/rent-the-runway-launches-rtr-platform-2018-11.

M. D. Ekstrand, A. Das, R. Burke, and F. Diaz. Fairness in recommender systems. In *Recommender Systems Handbook*, pages 679–707. 2022.

D. Fleder and K. Hosanagar. Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity. *Management science*, 55(5):697–712, 2009.

D. Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967.

R. Freeman, E. Micha, and N. Shah. Two-sided matching meets fair division. *arXiv preprint arXiv:2107.07404*, 2021.

D. Gale and L. S. Shapley. College admissions and the stability of marriage. *Americal Mathematical Monthly*, 69(1): 9–15, 1962.

S. Gollapudi, K. Kollias, and B. Plaut. Almost envy-free repeated matching in two-sided markets. In *International Conference on Web and Internet Economics*, pages 3–16, 2020.

A. Igarashi, Y. Kawase, W. Suksompong, and H. Sumita. Fair division with two-sided preferences. *arXiv preprint arXiv:2206.05879*, 2022.

A. S. Kelso and V. P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica: Journal of the Econometric Society*, pages 1483–1504, 1982.

Y. Koren. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 426–434, 2008.

R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 6th ACM Conference on Economics and Computation (EC)*, pages 125–131, 2004.

D. McFadden. Conditional logit analysis of qualitative choice behavior in frontier of econometrics. *Frontiers in Econometrics*, 1973.

J. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.

P. Nguyen, J. Dines, and J. Krasnodebski. A multi-objective learning to re-rank approach to optimize online marketplaces for multiple stakeholders. *arXiv preprint arXiv:1708.00651*, 2017.

T. T. Nguyen, M. Roos, and J. Rothe. A survey of approximability and inapproximability results for social welfare optimization in multiagent resource allocation. *Annals of Mathematics and Artificial Intelligence*, 68(1):65–90, 2013.

G. K. Patro, A. Biswas, N. Ganguly, K. P. Gummadi, and A. Chakraborty. Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms. In *Proceedings of The Web Conference 2020*, pages 1194–1204, 2020.

A. D. Procaccia. Cake cutting algorithms. In F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 13. Cambridge University Press, 2016.

M. Pycia. Stability and preference alignment in matching and coalition formation. *Econometrica*, 80(1):323–362, 2012.

S. Ramezani and U. Endriss. Nash social welfare in multiagent resource allocation. In *Proceedings of the 12th International Workshop on Agent-Mediated Electronic Commerce (AMEC)*, pages 117—131, 2010.

A. E. Roth. The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of political Economy*, 92(6):991–1016, 1984.

A. E. Roth. *Who gets what–and why: the new economics of matchmaking and market design*. Houghton Mifflin Harcourt, 2015.

A. E. Roth and M. Sotomayor. *Two-Sided Mathing: A Study in Game-Theoretic Modelling and Analysis*. Cambridge University Press, 1990.

A. E. Roth and M. Sotomayor. Two-sided matching. *Handbook of game theory with economic applications*, 1: 485–541, 1992.

J. M. Stalnaker. The matching program for intern placement. *Journal of Medical Education*, 28(11):13–19, 1953.

H. C. Williams. On the formation of travel demand models and economic evaluation measures of user benefit. *Environment and planning A*, 9(3):285–344, 1977.

D. Zeng and A. Psomas. Fairness-efficiency tradeoffs in dynamic fair division. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 911–912, 2020.

C. Ziegler, S. M. McNee, J. A. Konstan, and G. Lausen. Improving recommendation lists through topic diversification. In *Proceedings of the 14th international conference on World Wide Web*, pages 22–32, 2005.

A Proofs Omitted From The Main Body

An instance is called $\alpha-$balanced if, for all $i \in$, $u_{bi} \leq \lambda \cdot u_{ci}$ for all $b, c \in \mathcal{B}$.

An allocation A is called $\beta-$impartial when $u_b(A_b) \leq \beta \cdot u_c(A_C)$ for all $b, c \in \mathcal{B}$.

Consider 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 $A'_c = A_c \cup j \setminus i$.

Theorem 17. *Consider an α-balanced instance with β-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 upper bound as*

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

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

$$\begin{aligned}
\frac{\mathbb{P}(i, A')}{\mathbb{P}(i, A)} &= \frac{u_{bi}}{u_b(A_b) + u_{bi} - u_{bj}} \bigg/ \frac{u_{ci}}{u_c(A_c)} \\
&= \frac{u_{bi}}{u_{ci}} \cdot \frac{u_c(A_c)}{u_b(A_b) + u_{bi} - u_{bj}} \\
&\leq \frac{u_{bi}}{u_{ci}} \cdot \frac{u_c(A_c)}{u_b(A_b)} \\
&\leq \alpha\beta.
\end{aligned}$$

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

$$\begin{aligned}
&\mathbb{P}(i, A) < \mathbb{P}(i, A') \\
\iff \quad & \frac{u_{ci}}{u_c(A_c)} < \frac{u_{bi}}{b(A_b) + u_{bi} - u_{bj}} \\
\iff \quad u_b(A_b) + u_{bi} - u_{bj} &< \frac{u_{bi} \cdot u_c(A_c)}{u_{ci}}. \qquad (1)
\end{aligned}$$

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

$$\begin{aligned}
\frac{u_b(A'_b)}{u_b(A_b)} &= \frac{u_b(A_b) + u_{bi} - u_{bj}}{u_b(A_b)} \\
&< \frac{1}{u_b(A_b)} \cdot \frac{u_{bi} \cdot u_c(A_c)}{u_{ci}} && \text{(by Equation (1))} \\
&= \frac{u_{bi}}{u_{ci}} \cdot \frac{u_b(A_b)}{u_c(A_c)} \\
&\leq \alpha\beta
\end{aligned}$$

□

A.1 Identical Buyers

Theorem 7. *For buyers with identical preferences under unit exposure constraints, A^* is stable, PO and SEF1.*

We rely on the following technical lemma, which is straightforward to verify.

Lemma 8. *For all $x, y \geq 0$ satisfying $x + y = C$ for a fixed constant C, xy is monotonically increasing in $\min\{x, y\}$.*

Proof of Theorem 7. 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 b' be the buyer currently recommended i. By definition, a blocking pair implies $i \notin A_b^*$ and $\exists j \in A_b^*$ so that $u_{bj} < u_{bi}$ and $p(i, A^*) < p(i, A')$, where A' is constructed from A^* by exchanging i and j, i.e. $A'_b = A_b^* + i - j$ and $A'_{b'} = A^*_{b'} - i + j$ and $A^*_c = A'_c \, \forall c \notin \{b, b'\}$.

First, suppose $u(A_b^*) \geq u(A^*_{b'})$. It follows from $u_i > u_j$ that $u(A'_b) > u(A^*_b) \geq u(A^*_{b'}) > u(A'_{b'})$. But this implies $p(i, A^*) < p(i, A')$, contradicting that (b, i) is a blocking pair.

We may therefore assume that $u(A_b^*) < u(A^*_{b'})$. Since i increases their purchase probability by participating in the blocking pair,

$$\begin{aligned} \frac{u_i}{u(A^*_{b'})} < \frac{u_i}{u(A'_b)} = \frac{u_i}{u(A^*_b) + u_i - u_j} &\Leftrightarrow u(A^*_b) + u_i - u_j < u(A^*_{b'}) \\ &\Leftrightarrow u(A^*_b) + u_i < u(A^*_{b'}) + u_j. \end{aligned} \tag{2}$$

Finally, we compare $\min\{u(A^*_b), u(A^*_{b'})\}$ and $\min\{u(A'_b), u(A'_{b'})\}$. By assumption, $\min\{u(A^*_b), u(A^*_{b'})\} = u(A^*_b)$. It holds that $\min\{u(A'_b), u(A'_{b'})\} = \min\{u(A^*_b) + u_i - u_j, u(A^*_{b'}) + u_j - u_i\} > u(A^*_b)$, since $u_i - u_j > 0$ and $u(A^*_{b'}) + u_j - u_i > u(A^*_b) + u_i - u_j = u(A^*_b)$ by (2). It follows from Lemma 8 that $u(A'_b) \cdot u(A'_{b'}) > u(A^*_b) \cdot u(A^*_{b'})$, contradicting that A^* maximizes the product of buyer utilities with respect to the virtual values.

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

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 $b' \in \mathcal{B}$. By definition, $u(A^*_b) < u(A^*_{b'})$ and, for every $i \in A^*_b, j \in A^*_{b'}$, $u(A^*_b - i + j) < u(A^*_{b'} - j + i)$.

Suppose for contradiction that $u_i \geq u_j$ for all $i \in A^*_b, j \in A^*_{b'}$. Since $|A^*_b| = |A^*_{b'}|$, it follows that $u(A^*_b) \geq u(A^*_{b'})$, which is not the case. We conclude that there exists at least one pair of sellers (i, j) such that $u_i < u_j$. Let A' be the recommendation that results from swapping i and j, in other words, $A'_b = A^*_b + j - i$, $A'_{b'} = A^*_b - j + is$ and $A'_j = A^*_j$ for $j \in \mathcal{B} \setminus \{b, b'\}$.

Buyers have identical preferences and $A'_b \cup A'_{b'} = A^*_b \cup A^*_{b'}$, so $u(A'_b) + u(A'_{b'}) = u(A^*_b) + u(A^*_{b'})$. By the fact that A^* is not SEF1 and the choice of i and j, we observe

$$u(A^*_b) < u(A'_b) = u(A^*_b - i + j) < u(A^*_{b'} - j + i) = u(A'_{b'}) < u(A^*_{b'}),$$

however, this contradicts A^* maximizing the product of buyer virtual values by Lemma 8. We conclude that A^* is SEF1. ∎

Theorem 10. *Under general exposure constraints, A^* is PO and SEF1 for buyers with identical preferences.*

Proof. Pareto optimility 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 $S = \{(i, j) : i \in A^*_b \setminus X, j \in A^*_c \setminus X, u_i < u_j\}$ denote the set of feasible exchanges that are (strictly) improving for b. Suppose $S = \emptyset$. If $|X| = k$, then b and c are recommended idenitcal 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 $i^- = \text{argmin}_i\{u_i : i \in A_b^* \setminus X\}$ and $j^+ = \text{argmax}_j\{u_j : j \in A_c^* \setminus X\}$. If $S = \emptyset$, then in particular $(i^-, j^+) \notin S$ and, since this is a feasible exchange, it follows that $u_{i^-} > u_{j^+}$. Then

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

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

Select arbitrary $(i, j) \in S$. By assumption, $u(A_b^*) + u_j - u_i < u(A_c^*) + u_i - u_j$. 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 $d \in \mathcal{B} \setminus \{b, c\}$. Now

$$u(A_b^*) < u(A'_b) = u(A_b^*) + u_j - u_i < u(A_c^*) + u_i - u_j = u(A'_c) \leq u(A_c^*).$$

Since $u(A_b^*) + u(A_c^*) = u(A'_b) + u(A'_c)$, we conclude by Lemma 8 that $\prod_{b \in \mathcal{B}} u(A_b^*) < \prod_{b \in \mathcal{B}} u(A'_b)$, a contradiction. It follows that A^* is SEF1. □

A.2 Dichotomous Values

Recall that $v_{bi} \in \{a, a'\}$ for all $b \in \mathcal{B}$ and $i \in \mathcal{I}$, for some real-valued $a < a'$.

Theorem 13. *Under general exposure constraints, 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 contradition, 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, $u_{bi} = h$ and her choice set contains at least one low-valued item. Let $j \in A_b^*$ be such an item with $u_{bj} = \ell$. 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 $(u_{bi}, u_{cj'})$:

1. $(\ell, \cdot)$: Now $u_b(A'_b) \leq u_b(A_b^*)$. 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 $u_{bi} = \ell = u_{bj}$.

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

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

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

   Suppose instead $u_{ci} = h$. Now $u_c(A'_c) = u_c(A_c^*) + \ell - h \leq u_c(A_c^*)$. It follows that $u_c(A'_c) + u_b(A'_b) = u_c(A_c^*) + \ell - h + u_b(A_b^*) - \ell + h = u_c(A_c^*) + u_b(A_b^*)$, so we are in a position to check whether it is possible to apply Lemma 8. We know that $u_b(A_b^*) < u_b(A_b^*) + h - \ell = u_b(A'_b) < u_c(A_c^*)$, because both b and i benefit from participating in the blocing pair and $u_{bi} = h = u_{ci}$ by assumption. As a result, $u_c(A^*_c) - u_b(A^*_b) > h + \ell$. In contrast, $u_c(A^*_c) - u_c(A'_c) = h + \ell$. It follows that $u_b(A^*_b) < u_c(A'_c)$ and thus $\min\{u_b(A^*_b), u_c(A^*_c)\} < \min\{u_b(A'_b), u_c(A'_c)\}$. By Lemma 8, this contradicts A^* maximizing the product of virtual values.

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

A.3 Other Preference Domains

A.3.1 Latent factor models

First, we show that the instance of Theorem 3 (Table 2) can be factorised, which implies that stability can not be guaranteed when buyer values come from a latent factor model.

Example 18. *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.*

	a	b	c	d
1	2.34	0.08	2.04	1.82
2	2.23	2.06	1.38	1.49

	a	b	c	d
1	10.39	1.08	7.69	6.17
2	9.39	7.84	3.97	4.44

Table 8. Valuation matrix (left) and utility matrix (right).

It is straightforward to verify that this instance allows the same blocking pairs identified in Theorem 3.

A.3.2 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.

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

Table 9. Buyers' virtual values: identical preference orders do not guarantee stable recommendations existing.

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

Proposition 19. *For two buyers with identical preference orders and four items* $\{a, b, c, d\}$ *with* $= 2$ *under unit exposure constraints, at least one of* $\{\{a, d\}, \{b, c\}\}$ *and* $\{\{b, c\}, \{a, d\}\}$ *are stable.*

A.4 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 benefit from participating in a blocking pair is as small as possible.

There are some obstacles to overcome. Stability depends on sellers' 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 find 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_{bs} which takes value 1 exactly when $s \in \mathcal{I}$ 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

$$\sum_{b \in \mathcal{B}} x_{bs} = c_s, \forall s \in \mathcal{I}, \tag{3}$$

$$\sum_{s \in \mathcal{I}} x_{bs} = k, \forall b \in \mathcal{B}. \tag{4}$$

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

$$g \geq \frac{u_{bs}}{\sum_{k \in \mathcal{I}} x_{bk} u_{bk} - u_{bt} + u_{bs}} \Big/ \frac{u_{cs}}{\sum_{k \in \mathcal{I}} x_{ck} u_{ck}},$$

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

$$g \cdot (\sum_{k \in \mathcal{I}} x_{bk} u_{bk} - u_{bt} + u_{bs}) \geq \frac{u_{bs}}{u_{cs}} \cdot (\sum_{k \in \mathcal{I}} x_{ck} u_{ck}).$$

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

$$z_{bs} \geq 0, \tag{5}$$

$$z_{bs} \leq x_{bs} \cdot G, \tag{6}$$

$$z_{bs} \leq g, \tag{7}$$

$$z_{bs} \geq e + (x_{bs} - 1)G, \tag{8}$$

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

$$\sum_{k \in \mathcal{I}} z_{bk} u_{bk} - g u_{bt} + g u_{bs} \geq \frac{u_{bs}}{u_{cs}} \cdot \sum_{k \in \mathcal{I}} x_{ck} u_{ck}, \tag{9}$$

which should hold for all $b \neq c \in \mathcal{B}, s \neq t \in \mathcal{I}$ as long as $x_{bt} = 1 = x_{cs}$ and $u_{bs} > u_{bt}$. To enforce this we create a new indicator variable

$$\delta_{bt}^{cs} = \begin{cases} 1, & \text{when } x_{bt} = 1 = x_{cs} \text{ and } u_{bs} > u_{bt}, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Now we can rewrite eq. (9) as

$$\sum_{k \in \mathcal{I}} z_{bk} u_{bk} - g u_{bt} + g u_{bs} \geq \frac{u_{bs}}{u_{cs}} \cdot \sum_{k \in \mathcal{I}} x_{ck} u_{ck} - (1 - \delta_{bt}^{cs}) M. \tag{10}$$

The following two constraints ensure that δ_{bt}^{cs} takes on the value 1 when expected,

$$(1-\delta_{bt}^{cs})M' \geq u_{bt} - u_{bs} - (x_{bt} + x_{bs} - 2)M, \tag{11}$$
$$-\delta_{bt}^{cs} M' \leq u_{bt} - u_{bs} - (x_{bt} + x_{bs} - 2)M, \tag{12}$$

for $M' > 2M$. Observe that when $x_{bt} = 1 = x_{cs}$ and $u_{bs} > u_{bt}$, eq. (11) does not bind and eq. (12) becomes $-\delta_{bt}^{cs} M' < 0$, implying $\delta_{bt}^{cs} = 1$. When $x_{cs} + x_{ct} < 2$, eq. (11) becomes $(1-\delta_{bt}^{cs})M' \geq 0$, ensuring $\delta_{bt}^{cs} = 0$, while (12) does not bind. Similarly $u_{bs} \leq u_{bt}$, implies $\delta_{bt}^{cs} = 0$.

This yields the mixed integer program

$$\begin{aligned}
&\min g \\
\text{s.t.}\quad & \text{eqs. (3) to (4)} && \text{(assignment constraints)} \\
& \text{eqs. (5) to (8) } \forall b \in \mathcal{B}, s \in \mathcal{I} && \text{(linearization constraints)} \\
& \text{eqs. (10) to (12) } \forall b \neq c \in \mathcal{B}, s \neq t \in \mathcal{I} && \text{(stablity constraints)} \\
& g \geq 0,\ x \in \{0,1\}^{|\mathcal{B}|\times|\mathcal{I}|},\ z \in \{0,1\}^{|\mathcal{B}|\times|\mathcal{I}|} \\
& \delta_{bt}^{cs} \in \{0,1\}^{|\mathcal{B}|\times|\mathcal{I}|\times|\mathcal{B}|\times|\mathcal{I}|}
\end{aligned}$$

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

A.4.1 Scalability

We compare the computational cost of the integer program with the three recommendation strategies in Section 4 by training SVD++, a matrix factorization based collaborative filter, on the datasets described in Section 5.1. An instance is created setting $k = 3$ and randomly selecting $B \in \{2, 2^2, \ldots, 2^8\}$ 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 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.

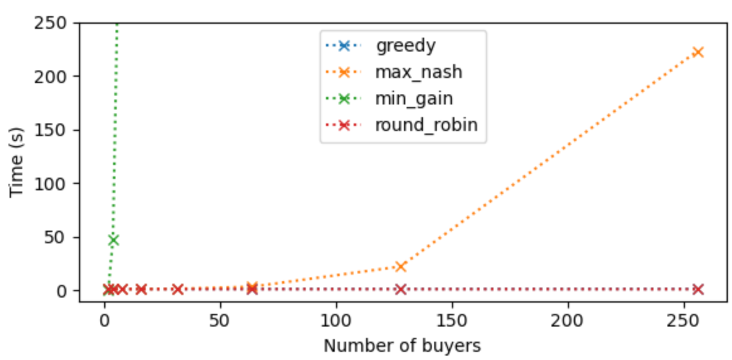

Figure 1. Time to find recommendations as function of the number of buyers with $k = 5$.