

# Stability in Online Assignment Games

Emile Martinez  
IRIT, Université Toulouse Capitole  
Toulouse, France  
emile.martinez@irit.fr

Felipe Garrido-Lucero  
IRIT, Université Toulouse Capitole  
Toulouse, France  
felipe.garrido-lucero@irit.fr

Umberto Grandi  
IRIT, Université Toulouse Capitole  
Toulouse, France  
umberto.grandi@irit.fr

## ABSTRACT

The assignment game models a housing market where buyers and sellers are matched, and transaction prices are set so that the resulting allocation is stable. Shapley and Shubik showed that every stable allocation is necessarily built on a maximum social welfare matching. In practice, however, stable allocations are rarely attainable, as matchings are often sub-optimal, particularly in online settings where agents arrive sequentially. In this paper, we leverage and compare two complementary measures of instability for allocations with sub-optimal matchings, establish their connections to the optimality ratio of the underlying matching, and use this framework to study the stability performances of randomized algorithms in online assignment games.

## KEYWORDS

Online Matching; Assignment Game; Optimality; Stability

### ACM Reference Format:

Emile Martinez, Felipe Garrido-Lucero, and Umberto Grandi. 2026. Stability in Online Assignment Games. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 9 pages. <https://doi.org/10.65109/YSRD5548>

## 1 INTRODUCTION

The assignment game of Shapley and Shubik [37] is a classical model of matching with transfers in which house buyers and house sellers pair together and set transaction prices. The objective is to find a stable allocation, i.e., a matching together with a price vector such that no pair of agents has an incentive to abandon their partners and trade with each other instead. Since its introduction, the model has been extensively studied and generalized [8, 10, 11, 14, 15, 22, 27, 36, 40].

Shapley and Shubik’s seminal work showed that stable allocations must be based on optimal matchings, i.e., matchings that maximize social welfare. However, in real-life scenarios optimal solutions are rarely attainable, for example in online markets [7] and online advertising [33], typically modeled as online matching markets where algorithms like *Greedy* [19, 28] and *Ranking* [1, 26] are commonly applied. As a result, the implemented solutions are often *unstable*. Yet, most existing stability metrics are binary: they simply determine whether a matching is stable or not. This motivates the need for non-binary measures of stability that can assess how close a sub-optimal matching is to being stable, and their posterior application to analyze the stability of online algorithms.

**Contributions.** The contributions of the article are as follows.

- We leverage three metrics to evaluate allocations arising from sub-optimal matchings: the optimality ratio; the stability index, derived from the subset instability of Jagadeesan et al. [25]; and the  $\kappa$ -approximate core [17, 35, 41] from cooperative game theory.
- We prove that, for any allocation, the optimality ratio upper bounds the stability index, which in turn upper bounds the parameter  $\kappa$  of the  $\kappa$ -approximate core. This result refines and strengthens the classical connection between stability and optimality established by Shapley and Shubik.
- We show that whenever prices can be set *a posteriori* (either outside online applications or in settings where matching and pricing can be decoupled), the stability index can be maximized to exactly match the optimality ratio of the underlying matching.
- We initiate the study of randomized algorithms in *online assignment games*, where either buyers or edges between buyers and sellers arrive sequentially. We evaluate stability at three levels: ex-post (worst-case realization), ex-ante (expected performance), and average (performance under expected utilities). While ex-ante and average guarantees coincide when analyzing social welfare (e.g., the competitive ratio of online algorithms), they diverge under our non-linear stability notions.
- Building on existing literature and new results, we establish lower bounds for the stability metrics at the three levels above for both vertex-arrival and edge-arrival models, and complement them with tight examples in several cases.

**Related work.** Our article builds on two main literatures: stable matching and online matching. In the former, given agents with (possibly endogenous) preferences, the goal is to form a matching (possibly with additional elements such as prices) such that no pair of agents would prefer each other over their assigned partners. In the latter, the focus is on designing algorithms that make decisions as the market evolves and that guarantee high social welfare, ideally independent of the specific instance.

Despite the extensive literature on stable matchings, most work on uncertainty has focused on settings without transfers. In economics, several authors have studied dynamically stable matchings [5, 12, 29] as well as markets with incomplete information about agents’ preferences [3, 31]. In machine learning, research often considers markets with unknown preferences and addresses them through regret minimization [2, 6, 9, 30]. More recently, Min et al. [34] and Jagadeesan et al. [25] studied the assignment game with unknown utility functions, introducing instability metrics as notions of regret and designing reinforcement learning algorithms that achieve sublinear regret bounds.

Online bipartite matching [13, 23, 33] is one of the most fundamental problems in the online algorithms literature. It dates back to the seminal work of Karp et al. [26], who introduced the *Ranking* algorithm and proved its optimality in the unweighted case. Since



This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems ([www.ifaamas.org](http://www.ifaamas.org)). <https://doi.org/10.65109/YSRD5548>


then, research on online matching has focused on designing increasingly competitive algorithms, that is, algorithms that achieve social welfare closer to that of the offline optimum that knows all arrivals in advance. Aggarwal et al. [1] extended this result by showing that *Ranking* is also optimal in the vertex-weighted case.

For the more general edge-weighted matching problem, Feldman et al. [18] introduced the free disposal assumption, noting that without it no randomized algorithm can achieve a constant competitive ratio. Under this assumption, *Greedy* attains a  $1/2$ -competitive guarantee. Fahrback et al. [16] were the first to surpass this long-standing  $1/2$ -barrier, a result subsequently improved by several works [4, 21, 39].

**Outline.** The rest of the article is organized as follows. Section 2 introduces the Shapley–Shubik assignment game, presenting its main techniques and results. Section 3 presents the stability index and the approximate core, relates them to the optimality ratio, and presents a simple pricing procedure that guarantees a  $1/2$  ratio for stability. Section 4 initiates the study of randomized algorithms for online assignment games in both edge-arrival and vertex-arrival models. This section defines our three levels of performance guarantees (ex-post, ex-ante, and average) and establishes lower bounds for each metric, with several tight examples. Section 5 concludes.

## 2 THE ASSIGNMENT GAME

The assignment game, in the classical notation of Shapley and Shubik [37], consists of a tuple  $\Gamma := (B, S, \mathbf{h}, \mathbf{c})$  where  $B$  and  $S$  are finite agents sets which we name buyers and sellers,  $\mathbf{h} := (h_{i,j})_{i \in B, j \in S}$ , where  $h_{i,j} > 0$  represents the valuation of buyer  $i$  for seller  $j$ 's house, and  $\mathbf{c} := (c_j)_{j \in S}$ , where  $c_j > 0$  is the valuation of seller  $j$  for her house. Figure 1 shows an assignment game example with three buyers and two sellers.



**Figure 1: An assignment game instance.** Buyers' valuations are denoted over the edges while sellers' valuations are denoted under their houses.

Throughout the article, we will use  $i$  to denote a typical buyer and  $j$  to denote a typical seller.

**Definition 2.1.** A **matching** is an injective function  $\mu : B \cup S \rightarrow B \cup S$  such that, (1)  $\mu \circ \mu = \text{Id}$ , (2) for any  $i \in B$ ,  $\mu(i) \in S \cup \{i\}$ , and (3) for any  $j \in S$ ,  $\mu(j) \in B \cup \{j\}$ . Whenever  $\mu(i) = j$ , for  $i \in B$  and  $j \in S$ , we say that  $i$  and  $j$  are matched, while for any agent  $k \in B \cup S$ , such that  $\mu(k) = k$ , we say that the agent is unmatched.

Given a matched pair of agents  $(i, j) \in B \times S$ , we will alternatively write  $i = \mu_j$ ,  $j = \mu_i$ , or  $(i, j) \in \mu$ . Finally, we denote  $\mathcal{M}(B, S)$  to the set of all matchings between  $B$  and  $S$ .

**Definition 2.2.** An **allocation** is a pair  $(\mu, \mathbf{p})$ , where  $\mu \in \mathcal{M}(B, S)$  is a matching and  $\mathbf{p} \in \mathbb{R}_+^S$  is a price vector.

Given an allocation  $(\mu, \mathbf{p})$  and  $j$  a matched seller,  $p_j \in \mathbf{p}$  represents the price that the buyer  $\mu_j$  paid for  $j$ 's house. By convention, we assume  $p_j = c_j$  whenever  $j$  is unmatched.

**Definition 2.3.** Given an allocation  $(\mu, \mathbf{p})$ ,  $i \in B$ , and  $j \in S$ , we define the **agents' utilities** as

$$u_i(\mu, \mathbf{p}) := \begin{cases} h_{i, \mu_i} - p_{\mu_i} & \text{if } \mu_i \neq i, \\ 0 & \text{if } \mu_i = i. \end{cases}$$

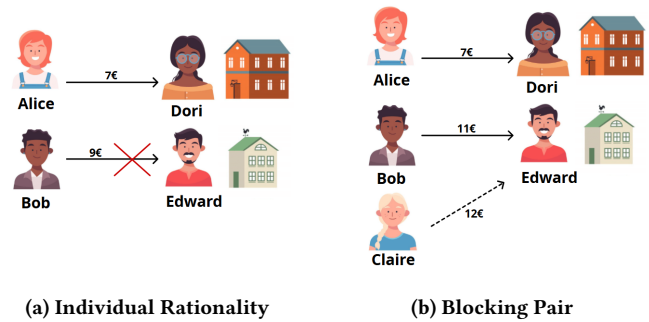
$$v_j(\mu, \mathbf{p}) := p_j - c_j.$$

For example, suppose that Alice and Dori in Figure 1 are matched and Alice pays 7€. Alice's utility is 3€ while Dori's utility is 1€.

**Definition 2.4.** An allocation  $(\mu, \mathbf{p})$  is called **stable** if it verifies

- $u_i(\mu, \mathbf{p}) \geq 0$  and  $v_j(\mu, \mathbf{p}) \geq 0$  for any  $i \in B$  and  $j \in S$ .
- There is no  $p \in \mathbb{R}_+$  and  $(i, j) \in B \times S$  such that  $h_{i,j} - p > u_i(\mu, \mathbf{p})$  and  $p - c_j > v_j(\mu, \mathbf{p})$ .

The first condition of Definition 2.4 corresponds to individual rationality, while the second one to the non-existence of blocking pairs. In our assignment game example, as illustrated in Figure 2a, matching Bob and Edward at price 9€ is not individually rational, as Edward prefers to be unmatched. Similarly, as illustrated in Figure 2b, matching Bob and Edward at price 11€ and letting Claire unmatched creates a blocking pair, Claire and Edward, as Claire can offer 12€ and strictly increase her and Edward's utility.



**Figure 2: Two possible allocations in our assignment game example.** The first allocation is not individually rational, while the second allocation has a blocking pair.

**Definition 2.5.** Let  $(\mu, \mathbf{p})$  be an allocation. We define its **social welfare** as

$$SW(\mu, \mathbf{p}) := \sum_{i \in B} u_i(\mu, \mathbf{p}) + \sum_{j \in S} v_j(\mu, \mathbf{p}).$$

It is easy to see that  $SW(\mu, \mathbf{p})$  does not depend on the price vector. In particular, we will alternatively denote it by  $SW(\mu)$ . A matching is called **optimal** if it maximizes the social welfare.

For any pair  $(i, j) \in B \times S$ , define  $a_{i,j} := h_{i,j} - c_j$ , which we refer to as their generated utility, and let  $\mathbf{a} := (a_{i,j})_{i \in B, j \in S}$ . We assume that  $a_{i,j} \geq 0$  for all  $(i, j) \in B \times S$ . This assumption is made without loss of generality thanks to the individual rationality property, as agents generating negative utility will prefer to remain unmatched rather than matching together.

**REMARK 1.** *The notion of generated utility arises from the cooperative game perspective of the assignment game. In this setting, an assignment game can be modeled as pairs of agents forming matches, with each match producing a surplus. This surplus, or generated utility, is then divided between the two parties, reflecting the transferable-utility nature of the model.*

Shapley and Shubik proved that stable allocations are necessarily based on an optimal matching. Optimal matchings can be obtained by solving a linear program, whose dual program outputs the respective agents utilities. The primal-dual linear programs considered by Shapley and Shubik are given below.

$$\begin{aligned}
 (P) \quad & \max \sum_{i \in B} \sum_{j \in S} a_{i,j} \cdot x_{i,j} \\
 \text{s.t.} \quad & \sum_{j \in S} x_{i,j} \leq 1, \forall i \in B, \\
 & \sum_{i \in B} x_{i,j} \leq 1, \forall j \in S, \\
 & x_{i,j} \in [0, 1], \forall (i, j) \in B \times S.
 \end{aligned}
 \quad
 \begin{aligned}
 (D) \quad & \min \sum_{i \in B} \alpha_i + \sum_{j \in S} \beta_j \\
 \text{s.t.} \quad & \alpha_i + \beta_j \geq a_{i,j}, \forall (i, j) \in B \times S, \\
 & \alpha_i, \beta_j \geq 0, \forall i \in B, \forall j \in S.
 \end{aligned}$$

Note that (P) always allows for integer optimal solutions thanks to the total-unimodularity of the matrix defined by the primal constraints. In particular, we will use interchangeably  $\mu$  and  $x$  when referring to matchings.

The utility vectors obtained as solutions of (D), by construction, verify all conditions in Definition 2.4. In terms of cooperative game theory, the utility vectors are said to belong to the core.

**Definition 2.6.** Given an assignment game  $\Gamma = (B, S, \mathbf{h}, \mathbf{c})$  and its corresponding matrix of generated utility  $\mathbf{a}$ , we define its **core** as,

$$C(\Gamma) := \{(u, v) \in \mathbb{R}_+^B \times \mathbb{R}_+^S \mid u_i + v_j \geq a_{i,j}, \forall (i, j) \in B \times S\}.$$

We state the main result of Shapley-Shubik's seminal article [37].

**THEOREM 2.7.** *Let  $\Gamma$  be an assignment game and  $(x, \alpha, \beta)$  be solutions of the pair primal-dual linear programs such that  $x$  is integral. It follows that  $(\alpha, \beta) \in C(\Gamma)$  and there exists a price vector  $\mathbf{p}$  such that  $(x, \mathbf{p})$  is a stable allocation, with  $u_i(x, \mathbf{p}) = \alpha_i$  and  $v_j(x, \mathbf{p}) = \beta_j$ , for any  $i \in B, j \in S$ .*

*Conversely, let  $(\mu, \mathbf{p})$  be a stable allocation. Then,  $\mu$  is an optimal matching and  $(u(\mu, \mathbf{p}), v(\mu, \mathbf{p})) \in C(\Gamma)$ .*

Note that the prices vector  $\mathbf{p}$  in the first part of Theorem 2.7's statement are trivially given by,

$$p_j := \beta_j + c_j, \text{ for any } j \in S,$$

as, by definition, the sellers' utility is given by the difference between the received payment and their valuation.

**Example 2.8.** In our running example, the optimal matching corresponds to matching Alice with Dori, Claire with Edward, and letting Bob unmatched. In addition, the optimal utilities can be

defined for example by Alice paying 6€ to Dori and Claire paying 12€ to Edward.

### 3 SUB-OPTIMALITY AND STABILITY

In real-world applications such as online markets, optimal matchings are rarely observed due to factors such as the sequential arrival of agents or the incomplete knowledge of their valuations. As stated by Shapley-Shubik [37] (cf. Theorem 2.7), no allocation can be stable if its underlying matching is sub-optimal. Hence, the purpose of this section is to measure the stability of sub-optimal matchings. To this end, we introduce two metrics: the stability index, based on the subset instability of Jagadeesan et al. [25], and the approximate core [17, 35, 41]. Each notion has its advantages and limitations. The stability index is comparatively easier to maximize, but it is sensitive to utility scaling and thus tends to focus on pairs that generate higher utility. The approximate core, while harder to achieve, avoids this bias and provides a more balanced measure of stability. We formalize this intuition by proving that for any allocation, the distance to the core is bounded by the stability index. Moreover, we prove that both stability measures are bounded by the optimality ratio of the underlying matching, showing that by improving stability we indirectly build optimal allocations (cf. Equation (3)).

First, we formalize the concept of sub-optimal matching.

**Definition 3.1.** Given an assignment game  $\Gamma$  and an allocation  $(\mu, \mathbf{p})$ , we define its **optimality ratio** as

$$\lambda(\mu, \mathbf{p}) = \lambda(\mu) := \frac{SW(\mu)}{OPT},$$

where OPT denotes the social welfare of an optimal matching, that is, the solution of the primal problem (P).

#### 3.1 Stability Index

The stability index is based on the concept of subset instability [25], originally introduced as a notion of regret, that measures, for every sub-coalition, the difference between the social welfare it actually obtained and the maximum it could have obtained.

**Definition 3.2.** We define the **stability index** of a given allocation  $(\mu, \mathbf{p})$  as

$$\mathcal{J}(\mu, \mathbf{p}) := 1 - \frac{\mathcal{I}(\mu, \mathbf{p})}{OPT},$$

with  $\mathcal{I}(\mu, \mathbf{p})$  the subset instability of  $(\mu, \mathbf{p})$ , formally defined as


$$\mathcal{I}(\mu, \mathbf{p}) := \max_{(B', S') \subseteq (B, S)} \max_{\mu' \in \mathcal{M}(B', S')} \{SW(\mu') - SW|_{B', S'}(\mu, \mathbf{p})\},$$

where  $\mu'$  matches only agents from  $B'$  to  $S'$  and  $SW|_{B', S'}(\mu, \mathbf{p})$  denotes the social welfare of  $(\mu, \mathbf{p})$  restricted to the sub-market  $(B', S')$ , that is,

$$SW|_{B', S'}(\mu, \mathbf{p}) := \sum_{i \in B'} u_i(\mu, \mathbf{p}) + \sum_{j \in S'} v_j(\mu, \mathbf{p}).$$

Note that subset instability does not require a price vector on the sub-market  $(B', S')$  as prices cancel each other when computing the social welfare  $SW(\mu')$ . However, since  $(\mu, \mathbf{p})$  is an allocation on the whole market  $(B, S)$ , the prices do not necessarily cancel out in  $SW|_{B', S'}(\mu, \mathbf{p})$ .

*Example 3.3.* Consider the allocation  $(\mu, \mathbf{p})$  illustrated in Figure 3. It follows that  $\lambda(\mu, \mathbf{p}) = \frac{2}{3}$  as both couples generate, respectively,  $4\epsilon$  and  $2\epsilon$ , while the optimal matching has social welfare equal to  $9\epsilon$ . Regarding the stability index, note that considering the submarket defined by Alice, Claire, Dori, and Edward, and matching Alice with Dori and Claire with Edward, the difference on social welfare is equal to  $4\epsilon$ . Since no other combination of submarket and matching creates a higher difference,  $\mathcal{I}(\mu, \mathbf{p}) = 4\epsilon$ , and, thus,  $\mathcal{J}(\mu, \mathbf{p}) = \frac{5}{9}$ .



**Figure 3: An allocation with  $\mathcal{J}(\mu, \mathbf{p}) = \frac{5}{9}$ .**

Jagadeesan et al. [25] showed that the subset instability of a given allocation is always at least as large as the additive optimality gap of its matching. Consequently, the stability index is always upper-bounded by the optimality ratio, as illustrated in Example 3.3. As a complementary result, we show that by carefully choosing the prices of an allocation, the stability index can be maximized to exactly match the optimality ratio. To this end, we consider the notion of a stabilizing subsidy and recall a technical result, both drawn from Jagadeesan et al. [25].

**Definition 3.4.** Let  $(\mu, \mathbf{p})$  be an allocation. We define the **minimum stabilizing subsidy** as the solution of the following problem,

$$\begin{aligned} \min_{(\tau, \eta) \in \mathbb{R}_+^B \times \mathbb{R}_+^S} \quad & \sum_{i \in B} \tau_i + \sum_{j \in S} \eta_j \\ \text{s.t.} \quad & u_i(\mu, \mathbf{p}) + \tau_i \geq 0, \quad \forall i \in B, \\ & v_j(\mu, \mathbf{p}) + \eta_j \geq 0, \quad \forall j \in S, \\ & u_i(\mu, \mathbf{p}) + \tau_i + v_j(\mu, \mathbf{p}) + \eta_j \geq a_{i,j}, \quad \forall (i, j) \in B \times S. \end{aligned}$$

A minimum stabilizing subsidy corresponds to the minimum utility injection required to make the allocation  $(\mu, \mathbf{p})$  stable.

**LEMMA 3.5.** *For any allocation, the minimum stabilizing subsidy is equal to the subset instability.*

We are now able to show the following result.

**THEOREM 3.6.** *For any allocation  $(\mu, \mathbf{p})$ , it always holds  $\mathcal{J}(\mu, \mathbf{p}) \leq \lambda(\mu)$ . Moreover, for any matching  $\mu$ , there exists  $\mathbf{p} \in \mathbb{R}_+^S$  such that,  $(\mu, \mathbf{p})$  is individually rational and  $\mathcal{J}(\mu, \mathbf{p}) = \lambda(\mu)$ .*

**PROOF.** The upper bound for the stability index is a consequence of the results of Jagadeesan et al. [25]. We focus then on proving the second part of the statement. As observed by Shi [38], without loss of generality, we can assume that  $|B| = |S|$  and that  $\mu$  is a maximum size matching, i.e., no agent in  $B \cup S$  is unmatched. By Lemma 3.5, minimizing the subset instability over the prices vector

can be written as

$$\begin{aligned} \min_{(\mathbf{p}, \tau, \eta) \in \mathbb{R}_+^S \times \mathbb{R}_+^B \times \mathbb{R}_+^S} \quad & \sum_{i \in B} \tau_i + \sum_{j \in S} \eta_j \\ \text{s.t.} \quad & h_{i, \mu_i} - p_{\mu_i} + \tau_i \geq 0, \quad \forall i \in B, \\ & p_j - c_j + \eta_j \geq 0, \quad \forall j \in S, \\ & h_{i, \mu_i} - p_{\mu_i} + \tau_i + p_j - c_j + \eta_j \geq a_{i,j}, \quad \forall (i, j) \in B \times S. \end{aligned}$$

The dual of the previous problem is (see [32]),

$$\begin{aligned} \max \quad & \sum_{i \in B} \sum_{j \in S} \gamma_{i,j} (h_{i,j} - h_{i, \mu_i}) + \sum_{j \in S} \beta_j c_j - \sum_{i \in B} \alpha_i h_{i, \mu_i} \\ \text{s.t.} \quad & \alpha_i + \sum_{j \in S} \gamma_{i,j} = 1, \quad \forall i \in B, \\ & \beta_j + \sum_{i \in B} \gamma_{i,j} \leq 1, \quad \forall j \in S, \\ & \alpha, \beta, \gamma \geq 0. \end{aligned}$$

Using the first constraint, note that the objective function becomes:

$$\sum_{i \in B} \sum_{j \in S} \gamma_{i,j} h_{i,j} + \sum_{j \in S} \beta_j c_j - \sum_{i \in B} h_{i, \mu_i}.$$

The coefficients of  $\beta$  in the objective function being positive, the value of the problem is not modified by replacing the second class of constraints by

$$\beta_j + \sum_{i \in B} \gamma_{i,j} = 1, \quad \forall j \in S.$$

Using this, the objective function becomes

$$\sum_{i \in B} \sum_{j \in S} \gamma_{i,j} (h_{i,j} - c_j) + \sum_{j \in S} c_j - \sum_{i \in B} h_{i, \mu_i}.$$

Since none of the variables within  $\alpha$  and  $\beta$  appear in the objective function, they correspond to slack variables. Furthermore, as all agents are matched, it holds,

$$\sum_{i \in B} h_{i, \mu_i} - \sum_{j \in S} c_j = \sum_{(i,j) \in \mu} h_{i,j} - c_j = \text{SW}(\mu, \mathbf{p}).$$

Thus, the problem is finally written as,

$$\begin{aligned} \max_{\mathbf{p}} \quad & \sum_{i \in B} \sum_{j \in S} \gamma_{i,j} (h_{i,j} - c_j) - \text{SW}(\mu) \\ \text{s.t.} \quad & \sum_{j \in S} \gamma_{i,j} \leq 1, \quad \forall i \in B, \\ & \sum_{i \in B} \gamma_{i,j} \leq 1, \quad \forall j \in S, \\ & \beta, \gamma \geq 0. \end{aligned}$$

The resulting problem corresponding to problem (P) shifted of  $-\text{SW}(\mu)$ , we obtain that,

$$\min_{\mathbf{p}} \mathcal{I}(\mu, \mathbf{p}) = \max_{\mu' \in \mathcal{M}(B, S)} \text{SW}(\mu') - \text{SW}(\mu) = \text{OPT} - \text{SW}(\mu).$$

Regarding the choice of prices such that the corresponding allocation is individually rational, please refer to [32].  $\square$

As shown by Theorem 3.6, whenever prices can be chosen after the matching has been computed, the stability index can be maximized to exactly match the optimality ratio. While this represents an important theoretical bound, in most of online matching applications, prices must be decided at the same time that couples are

matched, in particular, ignoring the future pairs to be created. The following result provides a simple method to ensure a  $1/2$ -guarantee in such cases.

**PROPOSITION 3.7.** *Let  $\mu$  be a matching. Consider the price vector  $\mathbf{p}^{half} := (p_j^{half})_{j \in S} \in \mathbb{R}_+^S$ , defined by,*

$$p_j^{half} := \begin{cases} c_j + \frac{1}{2} \cdot a_{i,j} & \text{if } \mu_j = i, \\ c_j & \text{if } \mu_j = j. \end{cases}$$

*It follows that  $\frac{1}{2} \cdot \lambda(\mu) \leq \mathcal{J}(\mu, \mathbf{p}^{half})$ .*

**PROOF.** Let  $\mu$  and  $\mathbf{p}^{half}$  be as stated. Denote  $\mathbf{p}^* \in \operatorname{argmin}_{\mathbf{p}} \mathcal{I}(\mu, \mathbf{p})$  such that  $(\mu, \mathbf{p}^*)$  is individually rational (Theorem 3.6).

For  $i \in B$ , notice that  $u_i(\mu, \mathbf{p}^{half}) = \frac{1}{2} \cdot a_{i,\mu_i}$  and, since  $(\mu, \mathbf{p}^*)$  is individually rational,  $u_i(\mu, \mathbf{p}^*) \in [0, a_{i,\mu_i}]$ . Thus,  $u_i(\mu, \mathbf{p}^{half}) \geq \frac{1}{2} u_i(\mu, \mathbf{p}^*)$ . Similarly, for  $j \in S$ , it also holds that  $v_j(\mu, \mathbf{p}^{half}) \geq \frac{1}{2} \cdot v_j(\mu, \mathbf{p}^*)$ . Given  $(B', S') \subseteq (B, S)$  and  $\mu' \in \mathcal{M}(B', S')$ , it follows,

$$\begin{aligned} SW(\mu') - SW|_{B',S'}(\mu, \mathbf{p}^{half}) &= SW(\mu') - \sum_{i \in B'} u_i(\mu, \mathbf{p}^{half}) \\ &- \sum_{j \in S'} v_j(\mu, \mathbf{p}^{half}) \leq SW(\mu') - \frac{1}{2} \cdot \sum_{i \in B'} u_i(\mu, \mathbf{p}^*) - \frac{1}{2} \cdot \sum_{j \in S'} v_j(\mu, \mathbf{p}^*) \\ &= SW(\mu') - \frac{1}{2} \cdot SW|_{B',S'}(\mu, \mathbf{p}^*) \\ &= SW(\mu') - SW|_{B',S'}(\mu, \mathbf{p}^*) + \frac{1}{2} \cdot SW|_{B',S'}(\mu, \mathbf{p}^*) \\ &\leq \mathcal{I}(\mu, \mathbf{p}^*) + \frac{1}{2} \cdot SW|_{B',S'}(\mu, \mathbf{p}^*) \\ &\leq \text{OPT} - SW(\mu) + \frac{1}{2} \cdot SW(\mu) = \text{OPT} - \frac{1}{2} \cdot SW(\mu), \end{aligned}$$

where the last inequality uses the fact that  $\mathcal{I}(\mu, \mathbf{p}^*) = \text{OPT} - SW(\mu)$ . We conclude by taking the maximum over all sub-markets and matchings, and normalizing by OPT.  $\square$

Observe that prices in Proposition 3.7 only depend on the generated utility of the seller and buyer, making it computable online settings where agents or edges arrive sequentially to the market.

### 3.2 Approximated Core

The stability index is not scale-invariant with respect to the agents' utility. For example, consider the market in Figure 1 and change  $h_{\text{Alice,Dori}}$  to  $10^{10}$ . Matching only Alice and Dori, with a price equal to 6€, achieves a stability index close to 1, while leaving unmatched the rest of the agents. To avoid this, we consider an alternative stability notion, known as the approximate core [17, 35, 41].

**Definition 3.8.** Given an assignment game  $\Gamma$ , its corresponding matrix of generated utility  $\mathbf{a}$ , and  $\kappa \in [0, 1]$ , we define the  $\kappa$ -approximate core as

$$C_\kappa(\Gamma) := \{(u, v) \in \mathbb{R}_+^B \times \mathbb{R}_+^S \mid u_i + v_j \geq \kappa \cdot a_{i,j}, \forall (i, j) \in B \times S\}.$$

We denote  $(\mu, \mathbf{p}) \in C_\kappa(\Gamma)$  whenever  $(u(\mu, \mathbf{p}), v(\mu, \mathbf{p}))$  belongs to  $C_\kappa(\Gamma)$ , and say that  $(\mu, \mathbf{p})$  is in the  $\kappa$ -approximate core.

Matching only Alice and Dori in the modified market with  $h_{\text{Alice,Dori}} = 10^{10}$ , even though it achieves a stability index close to 1, it belongs to the 0-approximate core. With this in mind, as a first

result, we prove that belonging to the  $\kappa$ -approximate core is indeed stronger than achieving a stability index of  $\kappa$ .

**PROPOSITION 3.9.** *Let  $(\mu, \mathbf{p})$  be an allocation in the  $\kappa$ -approximate core. Then,  $\kappa \leq \mathcal{J}(\mu, \mathbf{p})$ .*

**PROOF.** Let  $(\mu, \mathbf{p})$  be a  $\kappa$ -approximate core allocation,  $(B', S') \subseteq (B, S)$  a sub-market, and  $\mu' \in \mathcal{M}(B', S')$  a matching. It follows,

$$\begin{aligned} SW(\mu') - SW|_{B',S'}(\mu, \mathbf{p}) &= \sum_{(i,j) \in \mu'} a_{i,j} - \sum_{i \in B'} u_i(\mu, \mathbf{p}) - \sum_{j \in S'} v_j(\mu, \mathbf{p}) \\ &= \sum_{(i,j) \in \mu'} a_{i,j} - u_i(\mu, \mathbf{p}) - v_j(\mu, \mathbf{p}) - \sum_{(i,i) \in \mu'} u_i(\mu, \mathbf{p}) - \sum_{(j,j) \in \mu'} v_j(\mu, \mathbf{p}) \\ &\leq \sum_{(i,j) \in \mu'} (a_{i,j} - u_i(\mu, \mathbf{p}) - v_j(\mu, \mathbf{p})) \\ &\leq \sum_{(i,j) \in \mu'} a_{i,j} - \kappa \cdot a_{i,j} = (1 - \kappa)SW(\mu') \leq (1 - \kappa)\text{OPT}, \end{aligned}$$

where the first inequality comes from individual rationality (utilities are non-negative in the  $\kappa$ -approximate core) and the second one from  $(\mu, \mathbf{p})$  being in the  $\kappa$ -approximate core. We conclude by taking maximum over all sub-markets and all matchings.  $\square$

The  $\kappa$ -approximate core is a *local* stability notion, as it evaluates the social welfare of each pair relative to their generated utility. In contrast, subset instability is a *global* stability notion, since it considers the aggregated social welfare across all pairs. Intuitively, if an allocation is locally close to being stable everywhere, then it must also be globally stable (Proposition 3.9), whereas the converse does not necessarily hold, as we show next.

**PROPOSITION 3.10.** *For any  $\kappa \in [0, 1)$ , there exists an assignment game with an allocation  $(\mu, \mathbf{p})$  verifying  $\kappa \leq \mathcal{J}(\mu, \mathbf{p}) \leq \lambda(\mu)$ , such that for no constant  $\kappa' \in (0, \kappa]$ ,  $(\mu, \mathbf{p})$  is in the  $\kappa'$ -approximate core.*

**PROOF.** Let  $\kappa \in [0, 1)$  be a constant. Consider an assignment game with two buyers  $B = \{a, b\}$ , two sellers  $S = \{\alpha, \beta\}$ , and the following matrix of generated utility  $\mathbf{a}$

$$\mathbf{a} = \begin{array}{c|cc} & \alpha & \beta \\ \hline a & \kappa & 0 \\ \hline b & 0 & 1 - \kappa \end{array}$$

Consider  $(\mu, \mathbf{p})$  defined by  $\mu = \{(a, \alpha), (b, b), (\beta, \beta)\}$  and  $\mathbf{p} = (0, 0)$ , that is, only  $a$  and  $\alpha$  are matched and  $a$  pays 0 to  $\alpha$ . The allocation verifies  $\text{OPT} - SW(\mu, \mathbf{p}) \leq \mathcal{I}(\mu, \mathbf{p}) \leq (1 - \kappa)\text{OPT}$ , however,  $(\mu, \mathbf{p})$  is not in the  $\kappa'$ -approximate stable for any  $\kappa' > 0$ .  $\square$

Interestingly, there exists a connection between the  $\kappa$ -approximate core and a multiplicative version of subset instability.

**THEOREM 3.11.** *Given  $(\mu, \mathbf{p})$  an individually rational allocation, define,*

$$\kappa(\mu, \mathbf{p}) := \min_{(i,j) \in B \times S} \frac{1}{a_{i,j}} \cdot (u_i(\mu, \mathbf{p}) + v_j(\mu, \mathbf{p})). \quad (1)$$

*Then, it always holds that  $(\mu, \mathbf{p})$  is in the  $\kappa(\mu, \mathbf{p})$ -approximated core. In addition,*

$$\kappa(\mu, \mathbf{p}) = \min_{(B',S') \subseteq (B,S)} \min_{\mu' \in \mathcal{M}(B',S')} \frac{SW|_{B',S'}(\mu)}{SW(\mu')}. \quad (2)$$

PROOF. Let  $(\mu, \mathbf{p})$  be an individually rational allocation. Recall that  $(\mu, \mathbf{p})$  is in the  $\kappa$ -approximate core, for  $\kappa$  some constant, if

$$\begin{aligned} & \forall (i, j) \in B \times S, u_i(\mu, \mathbf{p}) + v_j(\mu, \mathbf{p}) \geq \kappa \cdot a_{i,j} \\ \iff & \forall (i, j) \in B \times S, \frac{u_i(\mu, \mathbf{p}) + v_j(\mu, \mathbf{p})}{a_{i,j}} \geq \kappa \\ \iff & \min_{(i,j) \in B \times S} \frac{1}{a_{i,j}} \cdot (u_i(\mu, \mathbf{p}) + v_j(\mu, \mathbf{p})) \geq \kappa. \end{aligned}$$

Therefore,  $(\mu, \mathbf{p})$  always belongs to the  $\kappa(\mu, \mathbf{p})$ -approximated core, for  $\kappa(\mu, \mathbf{p})$  as in Equation (1). We prove next that Equation (2) holds. Consider

$$R := \min_{(B', S') \subseteq (B, S)} \min_{\mu' \in \mathcal{M}(B', S')} \frac{SW|_{B', S'}(\mu)}{SW(\mu')}.$$

It directly follows that  $\kappa(\mu, \mathbf{p}) \geq R$  as  $R$  considers all sub-markets, in particular those with only one agent per side. Consider next  $(B', S') \subseteq (B, S)$  and  $\mu'$  a matching from  $B'$  to  $S'$ . Consider, without loss of generality, that  $|B'| = |S'|$  and all agents are matched at  $\mu'$  (indeed, removing any unmatched agent from the coalition does not affect  $SW(\mu')$  and does not decrease  $SW|_{B', S'}(\mu)$ , by individual rationality). It follows,

$$\begin{aligned} \frac{SW|_{B', S'}(\mu)}{SW(\mu')} &= \frac{\sum_{(i,j) \in \mu'} u_i(\mu, \mathbf{p}) + v_j(\mu, \mathbf{p})}{\sum_{(i,j) \in \mu'} a_{i,j}} \\ &\geq \min_{(i,j) \in \mu'} \frac{u_i(\mu, \mathbf{p}) + v_j(\mu, \mathbf{p})}{a_{i,j}} \geq \kappa(\mu, \mathbf{p}). \end{aligned}$$

$\kappa(\mu, \mathbf{p})$  not depending on  $(B', S')$  nor  $\mu'$ , we conclude  $\kappa(\mu, \mathbf{p}) \leq R$ .  $\square$

Theorem 3.11 shows that an allocation will be as unstable as its most unstable couple. In the proof of Proposition 3.10, for example, the value of  $\kappa(\mu, \mathbf{p})$  of the constructed allocation is equal to 0.

To conclude the section, putting together Proposition 3.9 and Theorem 3.11, we conclude that for any allocation  $(\mu, \mathbf{p})$ , it holds

$$\kappa(\mu, \mathbf{p}) \leq \mathcal{J}(\mu, \mathbf{p}) \leq \lambda(\mu). \quad (3)$$

Equation (3) is particularly significant when applied to online matching, as it suggests that algorithms focusing on obtaining good stability bounds will invariably obtain good optimality bounds.

## 4 ONLINE STABLE ALLOCATIONS

This section considers randomized algorithms to find stable allocations in online assignment games. After adapting our stability metrics to uncertain settings, we obtain systematic bounds on optimality and stability in two well-known models of online matching.

### 4.1 Stability Under Uncertainty

We consider two standard online matching frameworks: the *edge arrival* model (Figure 4a) and the *vertex arrival* model (Figure 4b). In the former, we start with a bipartite graph containing only vertices, with edges arriving one by one. Upon the arrival of an edge, the algorithm must irrevocably decide whether to accept it, specifying a price to be paid, or to reject it. In the latter, we start with a bipartite graph with vertices fixed on one side, while vertices on the other side (together with their incident edges) arrive sequentially. Upon the arrival of such a vertex (an agent), the algorithm must decide

whether to match it to an available partner (possibly none) and, if matched, at what price. Whenever randomization is allowed in these decisions, we refer to the algorithm as randomized.

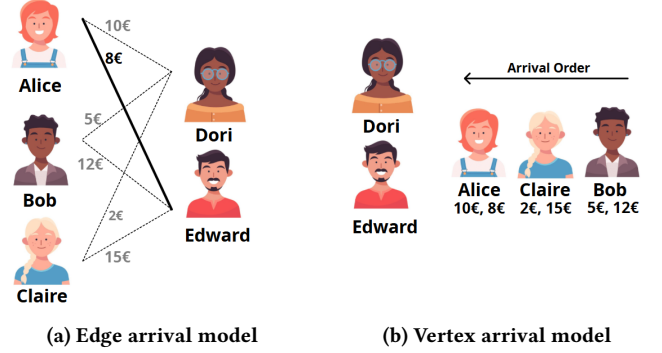


Figure 4: Two online matching markets models.

Given an online assignment game instance  $\Gamma = (B, S, \mathbf{h}, \mathbf{c})$  (either edge or vertex arrival)<sup>1</sup> and a randomized algorithm ALG, we denote  $\text{ALG}(\Gamma)$  the probability distribution of outcomes generated by ALG,  $\text{supp}(\text{ALG}(\Gamma))$  its support, that is, the set of possible outcomes of the algorithm on  $\Gamma$ , and  $(u(\text{ALG}, \Gamma), v(\text{ALG}, \Gamma))$  the vectors of expected utilities of the agents, that is, for any  $(i, j) \in B \times S$ ,

$$\begin{aligned} u_i(\text{ALG}, \Gamma) &= \mathbb{E}_{(\mu, \mathbf{p}) \sim \text{ALG}(\Gamma)} [u_i(\mu, \mathbf{p})], \\ v_j(\text{ALG}, \Gamma) &= \mathbb{E}_{(\mu, \mathbf{p}) \sim \text{ALG}(\Gamma)} [v_j(\mu, \mathbf{p})], \end{aligned}$$

where  $(\mu, \mathbf{p}) \sim \text{ALG}(\Gamma)$  indicates that allocations are sampled from the distribution induced by ALG in  $\Gamma$ . Given a metric  $m \in \{\lambda, \mathcal{J}, \kappa\}$ , we define the corresponding ex-post and ex-ante metric as follows,

$$\begin{aligned} m^{\text{post}}(\text{ALG}, \Gamma) &:= \min_{(\mu, \mathbf{p}) \in \text{supp}(\text{ALG}(\Gamma))} m(\mu, \mathbf{p}), \\ m^{\text{ante}}(\text{ALG}, \Gamma) &:= \mathbb{E}_{(\mu, \mathbf{p}) \sim \text{ALG}(\Gamma)} [m(\mu, \mathbf{p})]. \end{aligned}$$

Additionally, we define the average metric  $m^{\text{avg}}(\text{ALG}, \Gamma)$  by applying the definition of a metric  $m$  on the expected utilities of the agents, that is, on  $(u(\text{ALG}, \Gamma), v(\text{ALG}, \Gamma))$ .

REMARK 2. In both the online matching literature and the fair division literature, the ex-ante and average guarantees are usually treated as equivalent (see, e.g., [20, 24]). Indeed, whenever a metric is linear, its ex-ante and average version coincide. In particular, it always holds that  $\lambda^{\text{ante}}(\text{ALG}, \Gamma) = \lambda^{\text{avg}}(\text{ALG}, \Gamma)$ . However, the non-linearity of our stability metrics  $\mathcal{J}$  and  $\kappa$  breaks this equivalence, motivating the two different definitions in the randomized case (ex-ante and average), and showing that achieving stability in the online assignment game is a more subtle problem than achieving optimality.

We now show a useful result that allows us to systematize stability and optimality guarantees of randomized online algorithms:

PROPOSITION 4.1. Let  $\Gamma$  be an instance, ALG a randomized algorithm, and  $m \in \{\lambda, \mathcal{J}, \kappa\}$  a metric. It always holds,

$$m^{\text{post}}(\text{ALG}, \Gamma) \leq m^{\text{ante}}(\text{ALG}, \Gamma) \leq m^{\text{avg}}(\text{ALG}, \Gamma).$$

<sup>1</sup>For simplicity, we keep  $\Gamma$  to denote online assignment games.

Additionally, for any  $\gamma \in \{\text{post}, \text{ante}, \text{avg}\}$ , it always holds,

$$\kappa^\gamma(\text{ALG}, \Gamma) \leq \mathcal{J}^\gamma(\text{ALG}, \Gamma) \leq \lambda^\gamma(\text{ALG}, \Gamma).$$

The proof of the first part of Proposition 4.1 uses the linearity of the expected value and Jensen’s inequality. The second part adapts the arguments on the proof of Proposition 3.9. The formal proof is included in [32].

### 4.2 Equal Pricing and Edge Arrival Model

We begin by stating a generalization of Proposition 3.7 and by strengthening it to a tightness result that shows the impossibility of obtaining good stability guarantees under the edge arrival model.

**PROPOSITION 4.2.** *Let  $\Gamma$  be an instance and ALG a randomized matching algorithm. Denote ALG+Half the randomized algorithm that matches agents following ALG and, for any pair of matched agents  $(i, j) \in B \times S$ , the price  $p_j$  is set as in Proposition 3.7. For any  $\gamma \in \{\text{post}, \text{ante}, \text{avg}\}$ , it holds,*

$$\frac{1}{2} \cdot \lambda^\gamma(\text{ALG+Half}, \Gamma) \leq \mathcal{J}^\gamma(\text{ALG+Half}, \Gamma).$$

The proof of Proposition 4.2 follows similar arguments than Proposition 3.7 and can be found in [32].

The pricing algorithm *Half* considered in Proposition 4.2 may incorporate any randomized matching procedure and is applicable to both the edge and vertex arrival models. Interestingly, in the edge arrival setting we can construct two simple instances such that no algorithm from a broad family of randomized algorithms can achieve better than the  $1/2$ -factor on both instances simultaneously.

**PROPOSITION 4.3.** *Consider the two edge arrival instances  $\Gamma_1$  and  $\Gamma_2$  illustrated in Figure 5, where the first edge to arrive in each of them is between Alice and Dori, and all generated utilities are equal to 1.*




Figure 5: Two edge arrival instances

Let  $\mathcal{A}$  be the family of all randomized algorithms such that Alice and Dori are matched with probability 1. Then, for any  $\text{ALG} \in \mathcal{A}$  and any  $\gamma \in \{\text{post}, \text{ante}, \text{avg}\}$ , it holds,

$$\frac{1}{2} \cdot \lambda^\gamma(\text{ALG}, \Gamma_1) \geq \mathcal{J}^\gamma(\text{ALG}, \Gamma_1) \text{ or } \frac{1}{2} \cdot \lambda^\gamma(\text{ALG}, \Gamma_2) \geq \mathcal{J}^\gamma(\text{ALG}, \Gamma_2).$$

**PROOF.** From Proposition 4.1, given  $\text{ALG} \in \mathcal{A}$ , it is enough to prove that either  $\frac{1}{2} \cdot \lambda^{\text{post}}(\text{ALG}, \Gamma_1) \geq \mathcal{J}^{\text{avg}}(\text{ALG}, \Gamma_1)$  or  $\frac{1}{2} \cdot \lambda^{\text{post}}(\text{ALG}, \Gamma_2) \geq \mathcal{J}^{\text{avg}}(\text{ALG}, \Gamma_2)$ . Let  $\text{ALG} \in \mathcal{A}$  be a randomized algorithm. Suppose, once ALG matches Alice and Dori, it holds  $\mathbb{E}[u_{\text{Alice}}(\text{ALG})] \geq \mathbb{E}[v_{\text{Dori}}(\text{ALG})]$ . Notice that this assumption is independent on the chosen instance. Choose the first instance, where Bob was already present on the market and the second (and last) edge to arrive is between him and Dori. Since Dori is

already matched, the arriving edge is wasted. In this case, therefore,  $\lambda^{\text{post}}(\text{ALG}, \Gamma_1) = \frac{1}{2} = 1$  while

$$\mathcal{I}^{\text{avg}}(\text{ALG}, \Gamma_1) = a_{\text{Bob}, \text{Dori}} - \mathbb{E}[v_{\text{Dori}}(\text{ALG})] \geq \frac{1}{2},$$

since  $\mathbb{E}[u_{\text{Alice}}(\text{ALG})] + \mathbb{E}[v_{\text{Dori}}(\text{ALG})] = 1$ . We conclude,

$$\mathcal{J}^{\text{avg}}(\text{ALG}, \Gamma_1) = 1 - \frac{\mathcal{I}^{\text{avg}}(\text{ALG}, \Gamma_1)}{\text{OPT}} \leq \frac{1}{2} = \frac{1}{2} \cdot \lambda^{\text{post}}(\text{ALG}, \Gamma_1).$$

The proof if  $\mathbb{E}[u_{\text{Alice}}(\text{ALG})] \leq \mathbb{E}[v_{\text{Dori}}(\text{ALG})]$  is analogous considering the second edge arrival instance.  $\square$

Proposition 4.3 highlights the difficulty of obtaining good stability bounds in edge arrival models when studying stable matchings with transferable utility. As exposed by Rochford [36], achieving stability requires respecting the agents’ threat levels, i.e., the best utility each agent can secure outside their current match. In the edge arrival setting, however, these threat levels evolve dynamically on both sides of the market, which makes stability particularly challenging to maintain. This problem is only *partially* present in the vertex arrival model, as the threat levels of the buyers (the arriving agents) is much more determined at the moment of arrival.

### 4.3 Vertex Arrival Model

In this section, we distinguish between two variants of the vertex arrival model: vertex-weighted and edge-weighted with free disposal. In the vertex-weighted model, the generated utilities do not depend on the identity of the buyer; that is, for any  $(i, i', j) \in B \times B \times S$ , whenever both  $a_{i,j}$  and  $a_{i',j}$  are not zero, then  $a_{i,j} = a_{i',j}$ .

The edge-weighted model, by contrast, captures the classical situation in which different buyer–seller pairs may generate different utilities. In addition, the free-disposal assumption states that when a new buyer arrives, the decision-maker may unmatched a previously formed pair in order to reassign the seller to the arriving buyer, thereby leaving the previously matched buyer unmatched.

In each of these models, previous works have studied the design of *competitive* algorithms (see [13] for a complete survey). Informally, a randomized algorithm is competitive if it achieves a constant factor for the optimality ratio over all instances (the constant not depending on the instance). In our notation, an algorithm is competitive if  $\lambda^{\text{post}}$  (for deterministic algorithms) and  $\lambda^{\text{avg}}$  (for randomized algorithms) can be lower bounded by a constant not depending on the instance. Therefore, we can leverage the literature results to obtain stability guarantees for randomized algorithms.

**Vertex-weighted.** In the vertex-weighted setting, Karp et al. [26] introduced the *Ranking* algorithm which, in its alternative version (see Algorithm 1 in [32]), it randomly sets prices to goods before the arrival of any agent, and upon arrival of a buyer, it matches it greedily. Aggarwal et al. [1] proved that Ranking is optimal, that is, it achieves an average competitive ratio of  $1 - \frac{1}{e}$ . A closer examination of the proof of *Ranking*’s optimality shows that it defines transaction prices inducing expected utilities that belong to the  $\kappa$ -approximated core. In particular, we are able to restate the result of Aggarwal et al. [1] in our terminology:

**PROPOSITION 4.4.** *Consider the Ranking algorithm in the vertex-weighted setting. It holds,*

$$\min_{\text{Instance } \Gamma} \kappa^{\text{avg}}(\text{Ranking}, \Gamma) = 1 - \frac{1}{e}.$$

A more recent version of the proof of Ranking’s optimality (Theorem 5.6 in [13]) is based on a technical result (Lemma 5.5 in [13]), which corresponds exactly to Proposition 4.4. Plugging this result in Proposition 4.1, we conclude that Ranking achieves an average stability index and both ex-ante and average optimality ratio of  $1 - \frac{1}{e}$  over all assignment games.

Regarding  $\lambda^{\text{post}}$ , related to the competitive ratio of deterministic algorithms, Aggarwal et al. [1] proved that no algorithm can do better than  $1/2$  and that *Greedy*, the algorithm that matches the arriving vertex to the most profitable neighbor. We extend this bound to  $\kappa^{\text{post}}$  and  $\kappa^{\text{ante}}$  by complementing the Greedy algorithm with our pricing method *Half*, define in Proposition 4.2.

PROPOSITION 4.5. *Consider the vertex-weighted setting. It follows,*

$$\min_{\text{Instance } \Gamma} \lambda^{\text{post}}(\text{Greedy}, \Gamma) = \min_{\text{Instance } \Gamma} \kappa^\gamma(\text{Greedy} + \text{Half}, \Gamma) = \frac{1}{2},$$

with  $\gamma \in \{\text{post}, \text{ante}\}$  and moreover, no algorithm can do better in any of both cases.

PROOF SKETCH. The Greedy + *Half* algorithm sets prices so for any pair of agents  $(i, j) \in B \times S$ , each of them obtains at least half of utility they generate together, obtaining that the algorithm produces allocations in the  $\kappa$ -approximated core with  $\kappa \geq 1/2$ . The tightness then is proved by exhibiting an instance where the algorithm achieves a value  $\kappa$  exactly equal to  $1/2$ . The full proof is included in [32].  $\square$

Table 1 summarizes the our optimality and stability guarantees for the vertex-weighted setting.

**Table 1: Stability and optimality guarantees in vertex-weighted online assignment games. Remark all given values are tight, i.e., no algorithm can do better and for each of them, at least one algorithm achieves it. The results in bold correspond to our contributions, while the others are reformulations of literature results.**

	ex-post	ex-ante	avg
$\kappa$	<b>1/2</b>	<b>1/2</b>	$1 - 1/e$
$\mathcal{J}$	<b>1/2</b>	?	$1 - 1/e$
$\lambda$	1/2	$1 - 1/e$	$1 - 1/e$

Regarding the missing case, preliminary simulations suggest its value to be  $1 - 1/e$ .

**Edge-weighted with free disposal.** The free disposal assumption was introduced by Feldman et al. [18] due to the poor performances of online algorithms in the general edge-weighted model (Theorem 5.13 [13]). In our terminology, their result states that for any randomized matching algorithm ALG, it holds

$$\min_{\text{Instance } \Gamma} \lambda^{\text{avg}}(\text{ALG}, \Gamma) = 0.$$

In particular, from Proposition 4.1, it follows that no randomized algorithm can achieve a constant factor for any of our guarantees.

Under the free disposal assumption, in exchange, several works [4, 19, 28] have proved the following guarantees, which we reformulate in our notation:

PROPOSITION 4.6. *Consider the edge-weighted with free disposal setting. For any randomized algorithm ALG, it holds,*

$$\min_{\text{Instance } \Gamma} \lambda^{\text{post}}(\text{ALG}, \Gamma) \leq \frac{1}{2},$$

while, for any instance  $\Gamma$ ,  $\lambda^{\text{post}}(\text{Greedy}, \Gamma) \geq \frac{1}{2}$ . In addition, there exists a randomized algorithm ALG achieving,

$$\min_{\text{Instance } \Gamma} \lambda^{\text{ante}}(\text{ALG}, \Gamma) = \min_{\text{Instance } \Gamma} \lambda^{\text{avg}}(\text{ALG}, \Gamma) = 0.536.$$

Regarding the stability of these solutions, as for the edge arrival model, we obtain the following result.

PROPOSITION 4.7. *Consider the edge-weighted with free disposal setting. For any randomized algorithm ALG, it holds,*

$$\min_{\text{Instance } \Gamma} \kappa^{\text{post}}(\text{ALG}, \Gamma) = \min_{\text{Instance } \Gamma} \kappa^{\text{ante}}(\text{ALG}, \Gamma) = 0.$$

PROOF SKETCH. The proposition is proved by constructing two vertex-arrival instances, each of them with two buyers and two sellers. Similarly to the examples showed in Proposition 4.3, depending on the decision of the algorithm on the first arriving buyer, we choose the instance where the algorithm makes a mistake. In particular, it at least one of the two, the algorithm achieves an allocation in the 0-approximated core for  $\kappa^{\text{post}}$ . Regarding  $\kappa^{\text{ante}}$ , we construct an instance by considering several copies of the one created for  $\kappa^{\text{post}}$ . The full proof is included in [32].  $\square$

The previous results, combined with Proposition 4.2, are summarized in Table 2.

**Table 2: Stability and optimality guarantees in edge-weighted with free disposal online assignment games. Values with an inequality are lower bounds whose tightness remains open. Results in bold correspond to our contributions, while the others are reformulations of literature results.**

	ex-post	ex-ante	avg
$\kappa$	<b>0</b>	<b>0</b>	?
$\mathcal{J}$	$\geq 1/4$	$\geq \mathbf{0.268}$	$\geq \mathbf{0.268}$
$\lambda$	1/2	$\geq 0.536$	$\geq 0.536$

The value for  $\kappa^{\text{avg}}$  remains unknown, although we conjecture the value is zero. Interestingly, when allowing side payments, the results of Fahrbach et al. [16] can be used to show that  $\kappa^{\text{avg}} > \frac{1}{2}$ .

## 5 CONCLUSIONS

In this article, we initiated the study of stability in sub-optimal matchings and applied it to online assignment games, where either buyers or edges between buyers and sellers arrive sequentially. Our results show that stability naturally leads to optimality in the design of randomized algorithms for online matching, highlighting the study of stability in sub-optimal matchings as a promising and foundational research direction. As a direction for future work, a formal study of the dynamics governing the evolution of the agents’ bargaining power (their ability to influence the split of the generated utility) and threat levels (the best utility an agent can obtain outside of their assigned match) in online settings, paralleling the static analysis of Rochford [36], could provide valuable insights into the design of more stable allocation algorithms.

## ACKNOWLEDGMENTS

This work was supported by the ANR LabEx CIMI (grant ANR-11-LABX-0040) within the French State Programme “Investissements d’Avenir.”

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. This work is supported by ERC grant 101166894 “Advancing Digital Democratic Innovation” (ADDI).



Funded by  
the European Union



European Research Council  
Established by the European Commission

## REFERENCES

- [1] Gagan Aggarwal, Gagan Goel, Chinmay Karande, and Aranyak Mehta. 2011. Online vertex-weighted bipartite matching and single-bid budgeted allocations. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.
- [2] Soumya Basu, Karthik Abinav Sankararaman, and Abishek Sankararaman. 2021. Beyond  $\log^2(T)$  regret for decentralized bandits in matching markets. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- [3] Sushil Bikhchandani. 2017. Stability with one-sided incomplete information. *Journal of Economic Theory* 168 (2017).
- [4] Guy Blanc and Moses Charikar. 2022. Multiway online correlated selection. In *Proceedings of the 62nd Annual Symposium on Foundations of Computer Science (FOCS)*.
- [5] Martin Bullinger and René Romen. 2025. Stability in online coalition formation. *Journal of Artificial Intelligence Research* 82 (2025).
- [6] Sarah H Cen and Devavrat Shah. 2022. Regret, stability & fairness in matching markets with bandit learners. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- [7] Melvyn G Coles and Eric Smith. 1998. Marketplaces and matching. *International Economic Review* (1998), 239–254.
- [8] Vincent P Crawford and Elsie Marie Knoer. 1981. Job matching with heterogeneous firms and workers. *Econometrica: Journal of the Econometric Society* (1981).
- [9] Sanmay Das and Emir Kamenica. 2005. Two-Sided Bandits and the Dating Market. In *Proceedings of the International Joint Conferences on Artificial Intelligence (IJCAI)*, Vol. 5.
- [10] Gabrielle Demange and David Gale. 1985. The strategy structure of two-sided matching markets. *Econometrica: Journal of the Econometric Society* (1985).
- [11] Gabrielle Demange, David Gale, and Marilda Sotomayor. 1986. Multi-item auctions. *Journal of Political Economy* 94, 4 (1986).
- [12] Laura Doval. 2022. Dynamically stable matching. *Theoretical Economics* 17, 2 (2022).
- [13] Federico Echenique, Nicole Immorlica, and Vijay Vazirani. 2023. *Online and Matching-Based Market Design*.
- [14] Federico Echenique and Jorge Oviedo. 2004. Core many-to-one matchings by fixed-point methods. *Journal of Economic Theory* 115, 2 (2004).
- [15] Kimmo Eriksson and Johan Karlander. 2001. Stable outcomes of the roommate game with transferable utility. *International Journal of Game Theory* 29, 4 (2001).
- [16] Matthew Fahrbach, Zhiyi Huang, Runzhou Tao, and Morteza Zadimoghaddam. 2022. Edge-weighted online bipartite matching. *J. ACM* 69, 6 (2022).
- [17] Ulrich Faigle and Walter Kern. 1998. Approximate core allocation for binpacking games. *SIAM Journal on Discrete Mathematics* 11, 3 (1998).
- [18] Jon Feldman, Nitish Korula, Vahab Mirrokni, Shanmugavelayutham Muthukrishnan, and Martin Pál. 2009. Online ad assignment with free disposal. In *Proceedings of the International workshop on internet and network economics (WINE)*.
- [19] Marshall L Fisher, George L Nemhauser, and Laurence A Wolsey. 2009. An analysis of approximations for maximizing submodular set functions—II. In *Polyhedral Combinatorics: Dedicated to the memory of DR Fulkerson*.
- [20] Rupert Freeman, Nisarg Shah, and Rohit Vaish. 2020. Best of both worlds: Ex-ante and ex-post fairness in resource allocation. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*.
- [21] Ruiquan Gao, Zhongtian He, Zhiyi Huang, Zipei Nie, Bijun Yuan, and Yan Zhong. 2022. Improved online correlated selection. In *Proceedings of the 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*.
- [22] Felipe Garrido-Lucero and Rida Laraki. 2025. Stable matching games. *Social Choice and Welfare* (2025).
- [23] Zhiyi Huang, Zhihao Gavin Tang, and David Wajc. 2024. Online matching: A brief survey. *ACM SIGecom Exchanges* 22, 1 (2024).
- [24] Nicole Immorlica, Federico Echenique, and Vijay V Vazirani. 2023. *Online and Matching-Based Market Design*.
- [25] Meena Jagadeesan, Alexander Wei, Yixin Wang, Michael Jordan, and Jacob Steinhardt. 2021. Learning equilibria in matching markets from bandit feedback. *Advances in Neural Information Processing Systems (NeurIPS)* 34 (2021).
- [26] Richard M Karp, Umesh V Vazirani, and Vijay V Vazirani. 1990. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd annual ACM symposium on Theory of computing (STOC)*.
- [27] Alexander S Kelso Jr and Vincent P Crawford. 1982. Job matching, coalition formation, and gross substitutes. *Econometrica: Journal of the Econometric Society* (1982).
- [28] Benny Lehmann, Daniel Lehmann, and Noam Nisan. 2001. Combinatorial auctions with decreasing marginal utilities. In *Proceedings of the 3rd ACM conference on Electronic Commerce (EC)*.
- [29] Ce Liu. 2023. Stability in repeated matching markets. *Theoretical Economics* 18, 4 (2023).
- [30] Lydia T Liu, Horia Mania, and Michael Jordan. 2020. Competing bandits in matching markets. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- [31] Qingmin Liu, George J Mailath, Andrew Postlewaite, and Larry Samuelson. 2014. Stable matching with incomplete information. *Econometrica: Journal of the Econometric Society* 82, 2 (2014).
- [32] Emile Martinez, Felipe Garrido-Lucero, and Umberto Grandi. 2025. Stability in Online Assignment Games. *arXiv preprint arXiv:2510.09814* (2025).
- [33] Aranyak Mehta. 2013. Online Matching and Ad Allocation. *Foundations and Trends in Theoretical Computer Science* 8, 4 (2013).
- [34] Yifei Min, Tianhao Wang, Ruitu Xu, Zhaoran Wang, Michael Jordan, and Zhuoran Yang. 2022. Learn to match with no regret: Reinforcement learning in markov matching markets. *Advances in Neural Information Processing Systems (NeurIPS)* 35 (2022).
- [35] Xian Qiu and Walter Kern. 2016. Approximate core allocations and integrality gap for the bin packing game. *Theoretical Computer Science* 627 (2016).
- [36] Sharon C Rochford. 1984. Symmetrically pairwise-bargained allocations in an assignment market. *Journal of Economic Theory* 34, 2 (1984).
- [37] Lloyd S Shapley and Martin Shubik. 1971. The assignment game I: The core. *International Journal of game theory* 1, 1 (1971).
- [38] Peng Shi. 2023. Optimal matchmaking strategy in two-sided marketplaces. *Management Science* 69, 3 (2023).
- [39] Yongho Shin and Hyung-Chan An. 2021. Making Three out of Two: Three-Way Online Correlated Selection. In *32nd International Symposium on Algorithms and Computation (ISAAC 2021) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 212)*, Hee-Kap Ahn and Kunihiko Sadakane (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 49:1–49:17. <https://doi.org/10.4230/LIPIcs.ISAAC.2021.49>
- [40] Akiyoshi Shioura and others. 2017. On the Partnership formation problem. *Journal of Mechanism and Institution Design* 2 (2017).
- [41] Vijay V Vazirani. 2022. The general graph matching game: Approximate core. *Games and Economic Behavior* 132 (2022).